

Sixth Edition
Volume 2

Save 10%

on CompTIA® Exam
Vouchers

Coupon Inside!

CompTIA®

A+®

COMPLETE STUDY GUIDE

CORE 2 EXAM 220-1102

Includes one year of **FREE** access after activation to the
interactive online learning environment and study tools:

4 custom practice exams

200 electronic flashcards

Searchable key term glossary

**QUENTIN DOCTER
JON BUHAGIAR**

 **SYBEX**
A Wiley Brand

**Take the Next Step
in Your IT Career**

**Save
10%
on Exam Vouchers***
(up to a \$35 value)

*Some restrictions apply. See web page for details.

CompTIA®

Use coupon code WILEY10 during checkout. Redeeming the coupon code is easy:

1. Go to www.comptiastore.com.
2. Browse Certification Vouchers and select the exam voucher you want.
3. Add the voucher to the cart (note that for A+ you will need a separate voucher for each exam).
4. Enter the code WILEY10 on the purchase screen, click Apply and then click Proceed to Checkout to continue and complete the payment process.



CompTIA®

A+® Complete

Study Guide

Volume 2: Core 2 Exam 220-1202

Sixth Edition



Quentin Docter

Jon Buhagiar

 **SYBEX®**
A Wiley Brand

Copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada and the United Kingdom.

ISBNs for the 2-volume set: 9781394330034 (paperback), 9781394330058 (ePDF), 9781394330041 (ePub)

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permission.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: Wiley, the Wiley logo, and Sybex are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. CompTIA, A+, and the CompTIA logo are trademarks or registered trademarks of CompTIA, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and authors have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002. For product technical support, you can find answers to frequently asked questions or reach us via live chat at <https://sybexsupport.wiley.com>.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2025938071

Cover Design: Wiley

Cover Image: © Jeremy Woodhouse/Getty Images

For the fam.

—*Quentin Docter*

For my wife and son.

—*Jon Buhagiar*

Acknowledgments

As we were putting together this book, I was reminded of the proverb that begins “It takes a village . . .” That beginning definitely holds true for creating a book of this scope and size. From beginning to end, scores of dedicated professionals have focused on delivering the best book possible to you, the readers.

First, I need to thank my coauthor, Jon Buhagiar. I appreciate him diving in and dedicating himself to helping to produce an excellent book. I also need to give special thanks to our technical editor, Chris Crayton. He was meticulous and thorough, and he challenged me to always find new and better ways to communicate complex concepts. His entire focus was on providing the best training material possible, and I doubt there’s better in the business. Now, on to the rest of the team.

Kenyon Brown and Kim Wimpsett kept us on track and moving forward, which was a challenge at times. Saravanan Dakshinamurthy had the fun job of keeping us organized, which is akin to herding cats. Copyeditor Liz Welch reminded me yet again that I am no master of the English language and saved me from butchering it (too badly). Many thanks also go out to our proofreader, Nancy Carrasco, and our indexer, Tom Dinse. Without their great contributions, this book would not have made it into your hands.

On a personal note, I need to thank my family. My wife and kids are all incredibly supportive. Unfortunately, book writing as a side hustle while holding down a full-time job takes up a lot of time. I end up hiding in my office a lot, but they’re always there for me, and I couldn’t do it without them. Another huge thanks goes to my late grandpa, Joe, who got me into computers and taught me so many lessons I will never be able to repay. Finally, thanks to my friends who keep me relatively sane—Sean, Kurtis, Tim, John, Cory, Kevin, and others—and laugh at me when I tell them I spent my weekend writing about network cables and protocols.

—Quentin Docter

I would like to first thank my coauthor, Quentin Docter. Throughout the writing of this book, he helped me with his insight and his expertise in writing technical books. Without his words of wisdom and guidance, this book would not be the product it stands to be. I would also like to give special thanks to our technical editor, Chris Crayton. His thorough review of the material helped to identify many areas for us to elaborate on and polish.

I would also like to thank the many people who made this book possible: Kenyon Brown at Wiley Publishing, for giving me the opportunity to write this book and work with this wonderful team; Kim Wimpsett, for keeping us on track during the writing process; Pete Gaughan, who also kept us on track and organized during the publishing process; and our copyeditor, Liz Welch, for helping me use proper English. I’d also like to thank the many other people I’ve never met but who worked behind the scenes to make this book a success.

During the writing of this book, many others in my life supported me, and I'd like to take this opportunity to thank them as well. First and foremost, thanks to my wife and son for their support during the many evenings and weekends spent in my office at the computer. Finally, thanks to my friend and coworker, Bill, for encouraging me daily with his insight and jokes, as well as to all my other coworkers and friends. Thank you.

—Jon Buhagiar

About the Authors

Quentin Docter (A+, Network+, Tech+, Cloud Essentials +, MCSE, CCNA) is an IT author and consultant who started as a tech support agent for Packard Bell in 1994. Since then, he's worked in tech support, network administration, consulting, training, web development, and project management. During his career, he has achieved certifications from CompTIA, Microsoft, Cisco, Novell, and Sun Microsystems. He has written more than a dozen books for Sybex, including books on A+, Tech+, Cloud Essentials+, Server+, Windows, and Solaris 9 certifications, as well as PC hardware and maintenance.

Jon Buhagiar (Network+, A+, CCNA, MCSA, MCSE, BS/ITM) is an IT professional with two decades of experience in higher education. During the past 22 years he has been responsible for network operations at Pittsburgh Technical College and has led several projects, such as virtualization (server and desktop), VoIP, Microsoft 365, and many other projects supporting the quality of education at the college. He has achieved several certifications from Cisco, CompTIA, and Microsoft, and has taught many of the certification paths. He is the author of several books, including Sybex's *CompTIA Network+ Review Guide: Exam N10-008* (2021) and *CCNA Certification Practice Tests: Exam 200-301* (2020).

About the Technical Editor

Chris Crayton is a technical consultant, trainer, author and industry leading technical editor. He has worked as a computer technology and networking instructor, information security director, network administrator, network engineer, and PC specialist. Chris has authored several print and online books on PC repair, CompTIA A+, CompTIA Security+, and Microsoft Windows. He has also served as technical editor and content contributor on numerous technical titles for several of the leading publishing companies. He holds numerous industry certifications, has been recognized with many professional and teaching awards, and has served as a state-level SkillsUSA final competition judge.

Contents at a Glance

<i>Introduction</i>		<i>xxv</i>
<i>Assessment Test for Exam 220-1202</i>		<i>lxii</i>
<i>Answers to Assessment Test 220-1202</i>		<i>lxx</i>
Volume 2	CompTIA A+ Core 2 Exam (220-1202)	1
Chapter 1	Operating System Basics	3
Chapter 2	Windows Configuration	69
Chapter 3	Windows Administration	173
Chapter 4	Working with macOS and Linux	267
Chapter 5	Security Concepts	359
Chapter 6	Securing Operating Systems	453
Chapter 7	Troubleshooting Operating Systems and Security	539
Chapter 8	Scripting, Remote Access, and Artificial Intelligence	621
Chapter 9	Safety and Environmental Concerns	681
Chapter 10	Documentation and Professionalism	745
Appendix A	Answers to Review Questions	801
Appendix B	Answers to Performance-Based Questions	825
<i>Index</i>		<i>837</i>

Contents

<i>Introduction</i>	<i>xxv</i>
<i>Assessment Test for Exam 220-1202</i>	<i>lxii</i>
<i>Answers to Assessment Test 220-1202</i>	<i>lxx</i>

Volume 2	CompTIA A+ Core 2 Exam (220-1202)	1
-----------------	--	----------

Chapter 1	Operating System Basics	3
------------------	--------------------------------	----------

Understanding Operating Systems	6
Operating System Terms and Concepts	8
Operating Systems	10
Minimum System Requirements	15
Understanding Applications	17
System Requirements for Applications	18
Application Installation	20
Security Considerations	22
Other Considerations for New Applications	23
Cloud-Based Productivity Tools	23
Introduction to Windows	32
Windows Editions	33
Windows Features	35
The Windows Interface	42
What's in a Window?	53
File Management	57
Preparing for the Exam	61
Summary	61
Exam Essentials	62
Review Questions	64
Performance-Based Question	68

Chapter 2	Windows Configuration	69
------------------	------------------------------	-----------

Interacting with Operating Systems	72
Task Manager	72
Microsoft Management Console	77
Additional Tools	88
Control Panel	95
Windows Settings	127
The Windows Registry	155
Disk Management	158
Getting Disks Ready to Store Files and Programs	158
Checking the Health of Hard Disks and Optimizing Their Performance	166

	Summary	166
	Exam Essentials	167
	Review Questions	168
	Performance-Based Question	172
Chapter 3	Windows Administration	173
	Installing and Upgrading Windows	176
	Windows Installation Options	179
	The Installation Process	180
	The Upgrade Process	200
	Repair Installation	205
	Recovery Partition	206
	Side-by-Side Upgrading Windows 10/11	206
	Image Deployment	207
	Zero-Touch Deployment	207
	Upgrading Editions of Windows	209
	Security and Feature Updates	209
	Installation/Upgrade Boot Options	212
	Considerations	216
	Command-Line Tools	217
	Networking in Windows	236
	Networking Models	236
	User Authentication	243
	Establishing a Network Connection	244
	Accessing Resources	249
	Firewall Settings	252
	Client Network Configuration	256
	Summary	260
	Exam Essentials	260
	Review Questions	262
	Performance-Based Question	266
Chapter 4	Working with macOS and Linux	267
	macOS and Linux	271
	Applications on macOS	272
	Installing Applications from the App Store	272
	Installing Downloadable Applications	273
	Managing Applications	275
	Creating Shortcuts	276
	macOS Components	276
	Tools	277
	System Settings	279
	Features	289

macOS System Folders	295
<i>/Applications</i>	296
<i>/Users</i>	296
<i>/Library</i>	297
<i>/System</i>	298
Best Practices	299
Scheduled Backups	299
Scheduled Disk Maintenance	302
System Updates/App Store	302
Patch Management	304
Driver/Firmware Updates	305
Antivirus/Antimalware Updates	305
Rapid Security Response (RSR)	305
Linux Administration	306
Understanding the Syntax of Shell Commands	307
Discovering and Using Linux Commands	309
User Accounts	312
Maintaining the Operating System	315
Managing Processes	325
Directory Navigation	328
Directory Listings	330
Working with Files	334
Working with Directories	342
Networking Utilities	344
Getting Help	350
Summary	350
Exam Essentials	351
Review Questions	353
Performance-Based Question	357

Chapter 5	Security Concepts	359
	Physical Security Concepts	366
	Access Control Vestibule	366
	Badge Reader	367
	Video Surveillance	367
	Motion Sensors	368
	Alarm System	369
	Door Locks	369
	Equipment Locks	370
	Security Guards	372
	Fences	372
	Bollards	373
	Physical Security for Staff	374

Key Fobs	374
Smartcards and RFID Badges	375
Keys	375
Mobile Digital Key	376
Biometrics	376
Lighting	378
Magnetometers	378
Logical Security	379
Principle of Least Privilege	379
Zero Trust	379
Access Control Lists	380
Authentication Factors	382
Multifactor Authentication	382
Single Sign-on (SSO)	385
Just-in-Time Access	386
Privileged Access Management (PAM)	387
Mobile Device Management (MDM)	387
Data Loss Prevention	387
Identity and Access Management	388
Directory Services	388
Active Directory	389
Malware	396
Ransomware	396
Trojans	397
Keyloggers	397
Rootkits	398
Spyware	398
Cryptominers	398
Fileless	398
Stalkerware	399
Viruses	399
Virus Transmission in a Network	403
Botnets	404
Worms	404
Adware	404
Mitigating Software Threats	405
Antivirus	405
Antimalware	406
Detection and Response	407
Email Security Gateway Systems	408
Recovery Console	409
End-User Education	410
Software Firewalls	411
Reinstalling the OS	413

Social Engineering Attacks, Threats, and Vulnerabilities	414
Phishing	415
Shoulder Surfing	417
Tailgating	417
Impersonation	417
Dumpster Diving	418
Common Security Threats	418
Denial-of-Service Attacks	418
Evil Twin	420
Zero-Day Attacks	421
Spoofing Attacks	421
On-Path Attack	422
Password Attacks	422
Insider Threat	423
SQL Injection	424
Cross-Site Scripting (XSS)	424
Business Email Compromise (BEC)	425
Supply Chain/Pipeline Attacks	425
Exploits and Vulnerabilities	426
Noncompliant Systems	426
Patching and Updates	427
Operating Systems' Life Cycle	427
Unprotected Systems	428
BYOD	428
Security Best Practices	428
Data Encryption	429
Setting Strong Passwords	429
Requiring Passwords	431
Password Expiration	432
End-User Best Practices	432
Account Management	434
Disable AutoRun	439
Destruction and Disposal Methods	440
Recycling or Repurposing Best Practices	441
Physical Destruction	443
Summary	445
Exam Essentials	445
Review Questions	448
Performance-Based Question	452
Chapter 6 Securing Operating Systems	453
Working with Windows OS Security Settings	457
Users and Groups	457
User Authentication	460

NTFS vs. Share Permissions	466
Shared Files and Folders	474
System Files and Folders	477
Web Browser Security	488
Browser Download and Installation	489
Browser Patching	491
Extensions and Plug-Ins	491
Credentials Managers	493
Secure Data Transfers	494
Settings	495
Browser Features	502
Securing a SOHO Network (Wireless)	505
Changing Default Usernames and Passwords	505
Changing the SSID	506
Guest Network Isolation	506
Setting Encryption	507
Disabling SSID Broadcast	508
Wireless MAC Filtering	508
Frequencies	509
Radio Power Levels	511
Antenna and Access Point Placement	511
Assign Static IP Addresses	511
WPS	512
Authentication	512
Securing a SOHO Network (Wired)	513
Changing Default Credentials	514
Upgrading Firmware	514
Filtering	514
DHCP	518
Secure Management Access	519
Physical Security	519
Mobile Device Security	520
Hardening Techniques	520
Remote Wipes and Locator Applications	525
Remote Backup	528
Patch Management	528
Multifactor Authentication	529
Authenticator Applications	529
Policies and Procedures	530
IoT Considerations	531
Summary	532
Exam Essentials	532
Review Questions	534
Performance-Based Question	538

Chapter 7	Troubleshooting Operating Systems and Security	539
	Troubleshooting Common Microsoft Windows OS Problems	542
	Symptoms/Issues	543
	Troubleshooting Steps	567
	Troubleshooting Security Issues	580
	Common Symptoms	580
	Browser-Related Symptoms	585
	Best Practices for Malware Removal	589
	1. Investigate and Verify Malware Symptoms	589
	2. Quarantine Infected Systems	591
	3. Disable System Restore in Windows	591
	4. Remediate Infected Systems	592
	5. Schedule Scans and Run Updates	595
	6. Enable System Restore and Create a Restore Point in Windows	596
	7. Educate the End User	597
	Troubleshooting Mobile OS Issues	597
	Application Problems	597
	Performance Issues	600
	Random Reboots	601
	OS Fails to Update	602
	Extremely Short Battery Life	603
	Connectivity Issues	604
	Autorotate Issues	607
	Troubleshooting Mobile Security Issues	608
	Security Concerns	608
	Common Symptoms	611
	Summary	614
	Exam Essentials	615
	Review Questions	616
	Performance-Based Question	619
Chapter 8	Scripting, Remote Access, and Artificial Intelligence	621
	Scripting	623
	Scripting Basics	625
	Scripting Languages	629
	Scripting Use Cases	640
	Other Scripting Considerations	647
	Remote Access	649
	Remote Desktop Protocol	650
	Virtual Private Network (VPN)	651
	Virtual Network Computing	653
	Secure Shell	654
	Remote Monitoring and Management (RMM)	655

	Simple Protocol for Independent Computing Environments (SPICE)	656
	Windows Remote Management (WinRM)	657
	Remote Access Tools	657
	Security Considerations	659
	Artificial Intelligence	660
	Evolution	661
	Components	662
	Integration	669
	Limitations	670
	Considerations	673
	Summary	675
	Exam Essentials	675
	Review Questions	677
	Performance-Based Question	680
Chapter 9	Safety and Environmental Concerns	681
	Understanding Safety Procedures	684
	Identifying Potential Safety Hazards	684
	Creating a Safe Workplace	695
	Understanding Environmental Controls	706
	Managing the Physical Environment	706
	Handling and Disposing of Computer Equipment	713
	Understanding Policies, Licensing, and Privacy	722
	Dealing with Prohibited Content/Activity	723
	Incident Response	724
	Common Agreements	729
	Regulations and Compliance	730
	Managing Software Licenses	731
	Managing Sensitive Information	734
	Summary	737
	Exam Essentials	737
	Review Questions	739
	Performance-Based Question	743
Chapter 10	Documentation and Professionalism	745
	Documentation and Support	749
	Ticketing Systems	750
	Asset Management	755
	Common Documentation	758
	Knowledge Base/Articles	763
	Change Management Best Practices	765
	Documented Business Processes	765
	Backup and Recovery	773
	File-Level Backups	773
	Image-Based Backups	776

	Critical Applications	777
	Backup Media	777
	Backup Methods	779
	Data Restoration	780
	Backup Testing	780
	Demonstrating Professionalism	781
	Professional Appearance	781
	Communicating with Customers	782
	Using Appropriate Behavior	787
	Putting It All in Perspective	793
	Summary	794
	Exam Essentials	794
	Review Questions	796
	Performance-Based Question	800
Appendix A	Answers to Review Questions	801
	Chapter 1: Operating System Basics	802
	Chapter 2: Windows Configuration	803
	Chapter 3: Windows Administration	805
	Chapter 4: Working with macOS and Linux	808
	Chapter 5: Security Concepts	810
	Chapter 6: Securing Operating Systems	812
	Chapter 7: Troubleshooting Operating Systems and Security	815
	Chapter 8: Scripting, Remote Access, and Artificial Intelligence	817
	Chapter 9: Safety and Environmental Concerns	819
	Chapter 10: Documentation and Professionalism	821
Appendix B	Answers to Performance-Based Questions	825
	Chapter 1: Operating System Basics	826
	Chapter 2: Windows Configuration	826
	Chapter 3: Windows Administration	827
	Chapter 4: Working with macOS and Linux	831
	Chapter 5: Security Concepts	832
	Chapter 6: Securing Operating Systems	832
	Chapter 7: Troubleshooting Operating Systems and Security	833
	Chapter 8: Scripting, Remote Access, and Artificial Intelligence	834
	Chapter 9: Safety and Environmental Concerns	834
	Chapter 10: Documentation and Professionalism	835
	<i>Index</i>	837

Table of Exercises

Exercise	1.1	Changing a Screen Saver in Windows.	46
Exercise	1.2	Modifying the Taskbar	47
Exercise	1.3	Starting a Program from the Run Window	49
Exercise	2.1	Working with Task Manager	77
Exercise	2.2	Working with Performance Monitor.	86
Exercise	2.3	Changing the Time Zone	98
Exercise	2.4	Showing Hidden Files and Folders	103
Exercise	3.1	Command-Line Directory Management	219
Exercise	3.2	Running chkdsk within Windows	232
Exercise	3.3	Running chkdsk at the Command Line	232
Exercise	4.1	Installing Applications on macOS	274
Exercise	4.2	Uninstalling Applications on macOS	276
Exercise	4.3	Working with Files	341
Exercise	5.1	Testing Your Antimalware	407
Exercise	5.2	Testing Social Engineering	414
Exercise	6.1	Examining a SecurityToken	466
Exercise	6.2	Examining File Permissions	473
Exercise	6.3	Working with File Hashes	490
Exercise	6.4	Setting the Passcode Lock on an iPhone	521
Exercise	6.5	Setting the Passcode Lock on an Android Phone	522
Exercise	7.1	Reviewing Reliability Monitor.	556
Exercise	7.2	Manually Creating a Restore Point in Windows	574
Exercise	8.1	Creating and Running a Windows Batch Script	630
Exercise	8.2	Creating Your First PowerShell Script	634
Exercise	8.3	Creating a Script with AI	669
Exercise	9.1	Finding Trip Hazards	702
Exercise	9.2	Finding MSDSs	714
Exercise	10.1	Understanding Punctuality	787
Exercise	10.2	Surprise Someone.	791

Introduction

Welcome to the *CompTIA® A+® Complete Study Guide Volume 2: Core 2 Exam 220-1202 Sixth Edition*. This is the sixth edition of our best-selling study guide for the A+ certification sponsored by CompTIA (Computing Technology Industry Association). Thank you for choosing us to help you on your journey toward certification!

This book is written at an intermediate technical level; we assume that you already know how to *use* a personal computer and its basic peripherals, such as USB devices and printers, but we also recognize that you may be learning how to *service* some of that computer equipment for the first time. The exams cover basic computer service topics as well as more advanced issues, and they cover topics that anyone already working as a technician should be familiar with. The exams are designed to test you on these topics in order to certify that you have enough knowledge to fix and upgrade some of the most widely used types of personal computers and operating systems.

In addition to the information in the chapters, we've included a lot of extra material to help your study prep. At the end of each chapter is a list of exam essentials to know as well as 20 review questions to give you a taste of what it's like to take the exams. In addition, there are eight bonus exams of at least 75 questions each. Finally, there are flashcards designed to help your recall. Before you dive into those, though, we recommend you take the assessment test at the end of this introduction to gauge your current knowledge level.

Don't just study the questions and answers—the questions on the actual exams will be different from the practice ones included with this book. The exams are designed to test your knowledge of a concept or objective, so use this book to learn the objective behind the question. That said, we're confident that if you can do well on our quizzes, you will be well equipped to take the real exam.

This book covers more than just the exams, however. We believe in providing our students with a foundation of IT knowledge that will prepare them for real jobs, not just to pass a test. After all, life is not a multiple-choice test with the answers clearly laid out in front of you!

For experienced IT professionals, you can use the book to fill in the gaps in your current computer service knowledge. You may find, as many PC technicians have, that being well versed in all the technical aspects of hardware and operating systems is not enough to provide a satisfactory level of support—you must also have customer-relations skills, understand safety concepts, and be familiar with change management and environmental impacts and controls. We include helpful hints in all of these areas.

What Is CompTIA A+ Certification?

The A+ certification program was developed by CompTIA to provide an industrywide means of certifying the competency of computer service technicians. The A+ certification is granted to those who have attained the level of knowledge and troubleshooting skills that are needed

to provide capable support in the field of personal computers. It is similar to other certifications in the computer industry, such as the Cisco Certified Support Technician (CCST) program and the Microsoft certification programs. The theory behind these certifications is that if you need to have service performed on any of their products, you would sooner call a technician who has been certified in one of the appropriate certification programs than just call the first “expert” in the phone book.

The A+ certification program was created to offer a wide-ranging certification, in the sense that it is intended to certify competence with personal computers and mobile devices from many different makers/vendors. You must pass two tests to become A+ certified:

- The Core 1 220-1201 exam, which covers basic computer concepts, PC hardware, basic networking, mobile devices, and hardware troubleshooting
- The Core 2 220-1202 exam, which covers operating systems, security, software troubleshooting, and operational procedures

You don’t have to take the 220-1201 and the 220-1202 exams at the same time. However, the A+ certification is not awarded until you’ve passed both tests.

Why Become A+ Certified?

There are several good reasons to get your A+ certification. The CompTIA Candidate’s Information packet lists five major benefits:

- It demonstrates proof of professional achievement.
- It increases your marketability.
- It provides greater opportunity for advancement in your field.
- It is increasingly a requirement for some types of advanced training.
- It raises customer confidence in you and your company’s services.

Provides Proof of Professional Achievement The A+ certification is a status symbol in the computer service industry. Organizations that include members of the computer service industry recognize the benefits of A+ certification and push for their members to become certified. And more people every day are putting the “A+ Certified Technician” emblem on their business cards.

Increases Your Marketability A+ certification makes individuals more marketable to potential employers. A+ certified employees also may receive a higher base salary because employers won’t have to spend as much money on vendor-specific training.

Provides Opportunity for Advancement Most raises and advancements are based on performance. A+ certified employees work faster and more efficiently and are thus more productive. The more productive employees are, the more money they make for their company. And, of course, the more money they make for the company, the more valuable they are to the company. So, if an employee is A+ certified, their chances of being promoted are greater.

Fulfills Training Requirements Most major computer hardware vendors recognize A+ certification. Some of these vendors apply A+ certification toward prerequisites in their own respective certification programs, which has the side benefit of reducing training costs for employers.

Raises Customer Confidence As the A+ Certified Technician moniker becomes better known among computer owners, more of them will realize that the A+ technician is more qualified to work on their computer equipment than a noncertified technician.

How to Become A+ Certified

A+ certification is available to anyone who passes two exams. You don't have to work for any particular company. It's not a secret society. It is, however, an elite group. To become A+ certified, you must do two things:

- Pass the Core 1 220-1201 exam.
- Pass the Core 2 220-1202 exam.

The exams can be taken at any Pearson VUE testing center. If you pass both exams, you will get a certificate in the mail from CompTIA saying that you have passed, and you will also receive a lapel pin and business card.

To register for the tests, go to www.pearsonvue.com/comptia. You'll be asked for your name, Social Security number (an optional number may be assigned if you don't wish to provide your Social Security number), mailing address, phone number, employer, when and where you want to take the test, and your credit card number. (Payment arrangements must be made at the time of registration.)

Tips for Taking the A+ Exam

Here are some general tips for taking your exam successfully:

- Bring two forms of ID with you. One must be a photo ID, such as a driver's license. The other can be a major credit card or a passport. Both forms must include a signature.
- Arrive early at the exam center so that you can relax and review your study materials, particularly tables and lists of exam-related information. When you enter the testing room, you will need to leave everything outside; you won't be able to bring any materials into the testing area.
- Read the questions carefully. Don't be tempted to jump to an early conclusion. Make sure that you know exactly what each question is asking.
- Don't leave any unanswered questions. Unanswered questions are scored against you. Some questions will have multiple correct responses. When a question has more than one correct answer, a message at the bottom of the screen will prompt you to, for example, choose two. Be sure to read the messages displayed to know how many correct answers you must choose.

- When answering multiple-choice questions that you're not sure about, use a process of elimination to get rid of the obviously incorrect answers first. Doing so will improve your odds if you need to make an educated guess.
- For the hands-on questions where you need to complete a task, double-check that you have finished all parts of it, including checking all tabs available to you if they give you an interface to click on. Also, understand that these questions will take longer than multiple-choice ones. Spend the time needed on them; you will be able to make up time on the shorter multiple-choice questions.
- On form-based tests (nonadaptive), because the hard questions will take the most time, save them for last. You can move forward and backward through the exam.
- For the latest pricing on the exams and updates to the registration procedures, visit CompTIA's website at www.comptia.org.

Who Should Read This Book?

If you are one of the many people who want to pass the A+ exams, and pass them confidently, then you should buy this book and use it to study for the exams.

This book was written to prepare you for the challenges of the real IT world, not just to pass the A+ exams. This study guide will do that by describing in detail the concepts on which you'll be tested.

What Does This Book Cover?

CompTIA A+ Complete Study Guide: Exam 220-1201 and Exam 220-1202 covers everything you need to know to pass the CompTIA A+ exams.

Volume 2 of the book starts at Chapter 1 and concludes with Chapter 10. It covers all the topics on which you will be tested for Exam 220-1202:

Chapter 1: Operating System Basics Chapter 1 begins by covering the fundamentals of understanding various operating systems. The chapter then shifts to focus on supporting applications' installations. The chapter concludes by introducing Windows 10/11 editions, features, the Windows interface, and file management.

Chapter 2: Windows Configuration Chapter 2 explores Control Panel applets, Microsoft Management Console snap-ins, and Windows settings used to configure and customize the Windows operating system. The chapter also covers the Windows Registry and disk management.

Chapter 3: Windows Administration Chapter 3 guides you through the installation and upgrade of the Windows operating system. The chapter then explores the various command-line tools used to administer Windows. Finally, the chapter covers the administration and support of networking that is common with Windows.

Chapter 4: Working with macOS and Linux Chapter 4 covers common features and tools found in both the macOS and Linux operating systems. The chapter concludes by covering command-line tools used on Linux operating systems.

Chapter 5: Security Concepts Just when you think this book couldn't get any better, we toss in a chapter devoted to security. This chapter introduces key security concepts you need to know for the exam, including information on physical security, social engineering, wireless networks, best practices, and data destruction methods.

Chapter 6: Securing Operating Systems This chapter follows up on Chapter 5, which introduced security concepts, by providing specifics on how to secure Microsoft and mobile operating systems, prevent and remove malware, and configure network security on wired and wireless networks.

Chapter 7: Troubleshooting Operating Systems and Security Troubleshooting is a major part of almost every technician's job. In Chapter 7, we talk about troubleshooting common operating system issues as well as security challenges you may encounter.

Chapter 8: Scripting, Remote Access, and Artificial Intelligence Network administrators and technicians need to know how to create and modify scripts for automation. They also need remote access to systems in order to use or troubleshoot them. The chapter shows you technologies used to do just that. In addition, this chapter includes a wealth of knowledge on how artificial intelligence (AI) operates, as well as considerations for the adoption of AI.

Chapter 9: Safety and Environmental Concerns This chapter discusses the impacts computers can have on people as well as the environment. You learn how to avoid damaging a computer as well as avoid having it damage you, how to perform backups and disaster recovery, and how to properly dispose of computer components that are no longer needed.

Chapter 10: Documentation and Professionalism This chapter covers the "softer" side of working with computers. Specifically, you learn to practice proper communication, show professionalism, and implement best practices associated with documentation and change management.



Volume 1 of the book covers all the topics on which you will be tested for Exam 220-1201.

What's Included in the Book

We've included several learning tools throughout the book:

Assessment Tests At the end of this introduction is an assessment test that you can use to check your readiness for the exams. Take this before you start reading the book; it will help you determine the areas on which you might need to brush up. The answers

to the assessment test questions appear on a separate page after the last question of the test. Each answer includes an explanation and the chapter the material appears in.

Objective Map and Opening List of Objectives In this introduction, we have included a detailed exam objective map that shows you where each exam objective is covered. In addition, each chapter opens with a list of the exam objectives it covers. Use these resources to see exactly where each exam topic is covered.

Exam Essentials Each chapter, just after the summary, includes a number of “exam essentials.” These are the key topics that you should focus on when preparing for the exam.

Chapter Review Questions To test your knowledge as you progress through the book, each chapter ends with review questions. As you finish each chapter, answer the review questions and then check your answers—the correct answers and explanations are in Appendix A. You can go back to reread the section that deals with each question you got wrong to ensure that you answer correctly the next time that you’re tested on the material.

Performance-Based Questions The A+ exams may ask you to perform tasks to correctly answer a question. This may mean dragging and dropping items into the correct order or clicking the correct area of an image or screenshot. While it’s hard to simulate those in a paper book, the end of each chapter features a performance-based question. These questions will have you think through a topic presented in the chapter, and then perform a task. The correct answers are found in Appendix B.

Interactive Online Learning Environment and Test Bank

The interactive online learning environment that accompanies *CompTIA A+ Complete Study Guide: Exam 220-1201 and Exam 220-1202* provides a test bank with study tools to help you prepare for the certification exams and increase your chances of passing them the first time! The test bank includes the following elements:

Sample Tests All the questions and answers in this book are provided, including the assessment tests, which you’ll find at the end of this introduction, and the review questions at the end of each chapter. In addition, there are eight practice exams. Use these questions to test your knowledge of the study guide material. The online test bank runs on multiple devices.

Flashcards Four sets of questions are provided in digital flashcard format (a question followed by a single correct answer). You can use the flashcards to reinforce your learning and provide last-minute test prep before the exam.

Glossary The key terms from this book and their definitions are available as a fully searchable PDF.

eBooks Read eBook versions of this Study Guide in multiple formats on your favorite device.

Audio Instruction Over 2 hours of recorded instruction from the authors to help you study for the certification exam with confidence.



Go to www.wiley.com/go/sybextestprep to register and gain access to this interactive online learning environment and test bank with study tools.



Like all exams, the CompTIA A+ certification is updated periodically and may eventually be retired or replaced. At some point after CompTIA is no longer offering this exam, the old editions of our books and online tools will be retired. If you have purchased this book after the exam was retired, or are attempting to register in the Sybex online learning environment after the exam was retired, please know that we make no guarantees that this exam's online Sybex tools will be available once the exam is no longer available.

How to Use This Book

If you want a solid foundation for preparing for the A+ exams, this is the book for you. We've spent countless hours putting together this book with the intention of helping you prepare for the exams.

This book is loaded with valuable information, and you will get the most out of your study time if you understand how we put the book together. Here's a list that describes how to approach studying:

1. Take the assessment tests immediately following this introduction. It's okay if you don't know any of the answers—that's what this book is for. Carefully read over the explanations for any questions you get wrong, and make note of the chapters where that material is covered.
2. Study each chapter carefully, making sure you fully understand the information and the exam objectives listed at the beginning of the chapter. Again, pay extra-close attention to any chapter that includes material covered in questions you missed on the assessment test.
3. Read over the Summary and Exam Essentials. These will highlight the sections from the chapter you need to be familiar with before sitting for the exam.
4. Answer all the review questions at the end of each chapter. Specifically note any questions that confuse you, and study the corresponding sections of the book again. Don't just skim these questions. Make sure that you understand each answer completely.
5. Go over the electronic flashcards. These help you to prepare for the latest A+ exams, and they're really great study tools.
6. Take the practice exams.

Performance-Based Questions

CompTIA includes performance-based questions on the A+ exams. These are not the traditional multiple-choice questions with which you're probably familiar. These questions require the candidate to know how to perform a specific task or series of tasks. The candidate will be presented with a scenario and will be asked to complete a task. They will be taken to a simulated environment where they will have to perform a series of steps and will be graded on how well they complete the task.

The Sybex test engine does not include performance-based questions. However, at the end of each chapter, we have included a section called "Performance-Based Question," which is designed to measure how well you understood the chapter's topics. Some simply ask you to complete a task for which there is only one correct response. Others are more subjective, with multiple ways to complete the task. We provide the most logical or practical solution in Appendix B. Note that these questions may cover topic areas not covered in the actual A+ performance-based questions. However, we feel that being able to think logically is a great way to learn.

The CompTIA A+ Exam Objectives

The A+ exams consist of the Core 1 220-1201 exam and the Core 2 220-1202 exam. The following are the detailed exam objectives for Core 2 220-1202.

Exam objectives are subject to change at any time without prior notice and at CompTIA's sole discretion. Please visit the A+ Certification page of CompTIA's website (<https://comptia.org/certifications/a>) for the most current listing of exam objectives.

A+ Certification Exam Objectives: 220-1202

The following table lists the domains measured by this examination and the extent to which they are represented on the exam.

Domain	Percentage of exam
1.0 Operating Systems	28%
2.0 Security	28%
3.0 Software Troubleshooting	23%
4.0 Operational Procedures	21%
Total	100%

Objective Map

The following table lists where you can find the objectives covered in Volume 2 of this book.

Objective	Chapter(s)
1.0 Operating Systems	
1.1 Explain common operating system (OS) types and their purposes.	1, 2
▪ Workstation systems (OSs) ▪ Windows ▪ Linux ▪ macOS ▪ Chrome OS ▪ Mobile OSs ▪ iPadOS ▪ iOS ▪ Android ▪ Various filesystem types ▪ New Technology File System(NTFS) ▪ Resilient File System (ReFS) ▪ File Allocation Table 32 (FAT32) ▪ Fourth extended filesystem (ext4) ▪ Extended filesystem (XFS) ▪ Apple File System (APFS) ▪ Extensible File Allocation Table (exFAT) ▪ Vendor life-cycle limitations ▪ End-of-life (EOL) ▪ Update limitations ▪ Compatibility concerns between operating systems	

Objective	Chapter(s)
1.2 Given a scenario, perform OS installations and upgrades in a diverse environment. ▪ Boot methods ▪ Universal Serial Bus (USB) ▪ Network ▪ Solid-state/flash drives ▪ Internet-based ▪ External/hot-swappable drive ▪ Internal hard drive (partition) ▪ Multiboot ▪ Types of installations ▪ Clean install ▪ Upgrade ▪ Image deployment ▪ Remote network installation ▪ Zero-touch deployment ▪ Recovery partition ▪ Repair installation ▪ Other considerations ▪ Third-party drivers ▪ Partitioning ▪ GUID [globally unique identifier] partition Table (GPT) ▪ Master boot record (MBR) ▪ Drive format ▪ Upgrade considerations ▪ Backup files and user preferences ▪ Application and driver support/backward compatibility ▪ Hardware compatibility ▪ Feature updates ▪ Product life cycle	1, 2, 3

Objective	Chapter(s)
1.3 Compare and contrast basic features of Microsoft Windows editions.	1
▪ Windows 10 editions ▪ Home ▪ Pro ▪ Pro for Workstations ▪ Enterprise ▪ Windows 11 editions ▪ Home ▪ Pro ▪ Enterprise ▪ N versions ▪ Feature differences ▪ Domain vs. workgroup ▪ Desktop styles/user interface ▪ Availability of Remote Desktop Protocol (RDP) ▪ Random-access memory (RAM) support limitations ▪ BitLocker ▪ gpedit.msc ▪ Upgrade paths ▪ In-place upgrade ▪ Clean install ▪ Hardware requirements ▪ Trusted Platform Module (TPM) ▪ Unified Extensible Firmware Interface (UEFI)	

Objective	Chapter(s)
1.4 Given a scenario, use Microsoft Windows operating system features and tools.	2
▪ Task Manager ▪ Services ▪ Startup ▪ Performance ▪ Processes ▪ Users ▪ Microsoft Management Console (MMC) snap-in ▪ Event Viewer (eventvwr.msc) ▪ Disk Management (diskmgmt.msc) ▪ Task Scheduler (taskschd.msc) ▪ Device Manager (devmgmt.msc) ▪ Certificate Manager (certmgr.msc) ▪ Local User and Groups (lusrmgr.msc) ▪ Performance Monitor (perfmon.msc) ▪ Group Policy Editor (gpedit.msc) ▪ Additional tools ▪ System Information (msinfo32.exe) ▪ Resource Monitor (resmon.exe) ▪ System Configuration (msconfig.exe) ▪ Disk Cleanup (cleanmgr.exe) ▪ Disk Defragment (dfrgui.exe) ▪ Registry Editor (regedit.exe)	
1.5 Given a scenario, use the appropriate Microsoft command-line tools.	3
▪ Navigation ▪ cd ▪ dir ▪ Network ▪ ipconfig ▪ ping	

Objective	Chapter(s)
▪ netstat ▪ nslookup ▪ net use ▪ tracert ▪ pathping ▪ Disk management ▪ chkdsk ▪ format ▪ diskpart ▪ File management ▪ md ▪ rmdir ▪ robocopy ▪ Informational ▪ hostname ▪ net user ▪ winver ▪ whoami ▪ [command name] /? ▪ OS management ▪ gpupdate ▪ gpresult ▪ sfc	
1.6 Given a scenario, configure Microsoft Windows settings.	2
▪ Internet Options ▪ Devices and Printers ▪ Program and Features ▪ Network and Sharing Center ▪ System ▪ Windows Defender Firewall	

Objective	Chapter(s)
▪ Mail ▪ Sound ▪ User Accounts ▪ Device Manager ▪ Indexing Options ▪ Administrative Tools ▪ File Explorer Options ▪ View hidden files ▪ Hide extensions ▪ General options ▪ View options ▪ Power Options ▪ Hibernate ▪ Power plans ▪ Sleep/suspend ▪ Standby ▪ Choose what closing the lid does ▪ Turn on fast startup ▪ USB selective suspend ▪ Ease of Access ▪ Time and Language ▪ Update and Security ▪ Personalization ▪ Apps ▪ Privacy ▪ System ▪ Devices ▪ Network and Internet ▪ Gaming ▪ Accounts	

Objective	Chapter(s)
1.7 Given a scenario, configure Microsoft Windows networking features on a client/desktop. ▪ Domain joined vs. workgroup ▪ Shared resources ▪ Printers ▪ File servers ▪ Mapped drives ▪ Local OS firewall settings ▪ Application restrictions and exceptions ▪ Configuration ▪ Client network configuration ▪ Internet Protocol (IP) addressing scheme ▪ Domain Name System (DNS) settings ▪ Subnet mask ▪ Gateway ▪ Static vs. dynamic ▪ Establish network connections ▪ Virtual private network (VPN) ▪ Wireless ▪ Wired ▪ Wireless wide area network (WWAN)/cellular network ▪ Proxy settings ▪ Public network vs. private network ▪ File Explorer navigation–network paths ▪ Metered connections and limitations	3
1.8 Explain common features and tools of the macOS/desktop operating system. ▪ Installation and uninstallation of applications ▪ File types ▪ .dmg ▪ .pkg ▪ .app	4

Objective	Chapter(s)
App Store	
Uninstallation process	
▪ System folders ▪ /Applications ▪ /Users ▪ /Library ▪ /System ▪ /Users/Library ▪ Apple ID and corporate restrictions ▪ Best practices ▪ Backups ▪ Antivirus ▪ Updates/patches ▪ Rapid Security Response (RSR) ▪ System Preferences ▪ Displays ▪ Networks ▪ Printers ▪ Scanners ▪ Privacy ▪ Accessibility ▪ Time Machine ▪ Features ▪ Multiple desktops ▪ Mission Control ▪ Keychain ▪ Spotlight ▪ iCloud ▪ iMessage ▪ FaceTime ▪ Drive	

Objective	Chapter(s)
▪ Gestures ▪ Finder ▪ Dock ▪ Continuity ▪ Disk Utility ▪ FileVault ▪ Terminal ▪ Force Quit	
1.9 Identify common features and tools of the Linux client/desktop operating system.	4
▪ File management ▪ ls ▪ pwd ▪ mv ▪ cp ▪ rm ▪ chmod ▪ chown ▪ grep ▪ find ▪ Filesystem management ▪ fsck ▪ mount ▪ Administrative ▪ su ▪ sudo ▪ Package management ▪ apt ▪ dnf ▪ Network ▪ ip ▪ ping ▪ curl ▪ dig ▪ traceroute	

Objective	Chapter(s)
▪ Informational ▪ man ▪ cat ▪ top ▪ ps ▪ du ▪ df ▪ Text editors ▪ nano ▪ Common configuration files ▪ /etc/passwd ▪ /etc/shadow ▪ /etc/hosts ▪ /etc/fstab ▪ /etc/resolv.conf ▪ OS components ▪ systemd ▪ kernel ▪ bootloader ▪ Root account	
1.10 Given a scenario, install applications according to requirements.	1
▪ System requirements for applications ▪ 32-bit vs. 64-bit dependent application requirements ▪ Dedicated vs. integrated graphics card ▪ Video random-access memory (VRAM) requirements ▪ RAM requirements ▪ Central processing unit (CPU) requirements ▪ External hardware tokens ▪ Storage requirements	

Objective	Chapter(s)
▪ Application to OS compatibility ▪ Distribution methods ▪ Physical media vs. mountable ISO file ▪ Downloadable package ▪ Image deployment ▪ Impact considerations for new applications ▪ Device ▪ Network ▪ Operation ▪ Business	
1.11 Given a scenario, install and configure cloud-based productivity tools.	1
▪ Email systems ▪ Storage ▪ Sync/folder settings ▪ Collaboration tools ▪ Spreadsheets ▪ Videoconferencing ▪ Presentation tools ▪ Word processing tools ▪ Instant messaging ▪ Identity synchronization ▪ Licensing assignment	
2.0 Security	
2.1 Summarize various security measures and their purposes.	5
▪ Physical security ▪ Bollards ▪ Access control vestibule ▪ Badge reader ▪ Video surveillance ▪ Alarm systems ▪ Motion sensors	

Objective	Chapter(s)
▪ Door locks ▪ Equipment locks ▪ Security guards ▪ Fences ▪ Physical access security ▪ Key fobs ▪ Smart cards ▪ Mobile digital key ▪ Keys ▪ Biometrics ▪ Retina scanner ▪ Fingerprint scanner ▪ Palm print scanner ▪ Facial recognition technology (FRT) ▪ Voice recognition technology ▪ Lighting ▪ Magnetometers ▪ Logical security ▪ Principle of least privilege ▪ Zero Trust model ▪ Access control lists (ACLs) ▪ Multifactor authentication (MFA) ▪ Email ▪ Hardware token ▪ Authenticator application ▪ Short Message Service (SMS) ▪ Voice call ▪ Time-based one-time pass word (TOTP) ▪ One-time password/passcode (OTP) ▪ Security Assertions Markup Language (SAML) ▪ Single sign-on (SSO) ▪ Just-in-time access	

Objective	Chapter(s)
▪ Privileged access management (PAM) ▪ Mobile device management (MDM) ▪ Data loss prevention (DLP) ▪ Identity access management (IAM) ▪ Directory services	
2.2 Given a scenario, configure and apply basic Microsoft Windows OS security settings.	5, 6
▪ Defender Antivirus ▪ Activate/deactivate ▪ Update definitions ▪ Firewall ▪ Activate/deactivate ▪ Port security ▪ Application security ▪ User and groups ▪ Local vs. Microsoft account ▪ Standard account ▪ Administrator ▪ Guest user ▪ Power user ▪ Log-in OS options ▪ Username and password ▪ Personal identification number (PIN) ▪ Fingerprint ▪ Facial recognition ▪ SSO ▪ Passwordless/Windows Hello ▪ NTFS vs. share permissions ▪ File and folder attributes ▪ Inheritance ▪ Run as administrator vs. standard user ▪ User Account Control (UAC)	

Objective	Chapter(s)
▪ BitLocker ▪ BitLocker-To-Go ▪ Encrypting File System (EFS) ▪ Active Directory ▪ Joining domain ▪ Assigning log-in script ▪ Moving objects within organizational units ▪ Assigning home folders ▪ Applying Group Policy ▪ Selecting security groups ▪ Configuring folder redirection	
2.3 Compare and contrast wireless security protocols and authentication methods.	6
▪ Protocols and encryption ▪ Wi-Fi Protected Access 2 (WPA2) ▪ WPA3 ▪ Temporal Key Integrity Protocol (TKIP) ▪ Advanced Encryption Standard (AES) ▪ Authentication ▪ Remote Authentication Dial-in User Service (RADIUS) ▪ Terminal Access Controller Access-control System (TACACS+) ▪ Kerberos ▪ Multifactor	
2.4 Summarize types of malware and tools/methods for detection, removal, and prevention.	5
▪ Malware ▪ Trojan ▪ Rootkit ▪ Virus ▪ Spyware ▪ Ransomware	

Objective	Chapter(s)
▪ Keylogger ▪ Boot sector virus ▪ Cryptominer ▪ Stalkerware ▪ Fileless ▪ Adware ▪ Potentially unwanted program (PUP) ▪ Tools and methods ▪ Recovery console ▪ Endpoint detection and response (EDR) ▪ Managed detection and response (MDR) ▪ Extended detection and response (XDR) ▪ Antivirus ▪ Antimalware ▪ Email security gateway ▪ Software firewalls ▪ User education regarding common threats ▪ Antiphishing training ▪ OS reinstallation	
2.5 Compare and contrast common social engineering attacks, threats, and vulnerabilities.	5
▪ Social engineering ▪ Phishing ▪ Vishing ▪ Smishing ▪ QR code phishing ▪ Spear phishing ▪ Whaling ▪ Shoulder surfing ▪ Tailgating ▪ Impersonation ▪ Dumpster diving	

Objective	Chapter(s)
▪ Threats ▪ Denial of service (DoS) ▪ Distributed denial of service (DDoS) ▪ Evil twin ▪ Zero-day attack ▪ Spoofing ▪ On-path attack ▪ Brute-force attack ▪ Dictionary attack ▪ Insider threat ▪ Structured Query Language (SQL) injection ▪ Cross-site scripting (XSS) ▪ Business email compromise (BEC) ▪ Supply chain/pipeline attack ▪ Vulnerabilities ▪ Non-compliant systems ▪ Unpatched systems ▪ Unprotected systems (missing antivirus/missing firewall) ▪ EOL ▪ Bring your own device (BYOD)	
2.6 Given a scenario, implement procedures for basic small office/home office (SOHO) malware removal.	7
▪ 1. Investigate and verify malware symptoms. ▪ 2. Quarantine infected system. ▪ 3. Disable System Restore in Windows Home. ▪ 4. Remediate infected systems. a. Update antimalware software. b. Scan and removal techniques (e.g., safe mode, preinstallation environment) c. Reimage/reinstall. ▪ 5. Schedule scans and run updates. ▪ 6. Enable System Restore and create a restore point in Windows Home. ▪ 7. Educate the end user.	

Objective	Chapter(s)
2.7 Given a scenario, apply workstation security options and hardening techniques. ▪ Data-at-rest encryption ▪ Password considerations ▪ Length ▪ Character types ▪ Uniqueness ▪ Complexity ▪ Expiration ▪ Basic input/output system (BIOS)/Unified Extensible Firmware Interface (UEFI) passwords ▪ End-user best practices ▪ Use screensaver locks ▪ Log off when not in use ▪ Secure/protect critical hardware (e.g., laptops) ▪ Secure personally identifiable information (PII) and passwords ▪ Use password managers ▪ Account management ▪ Restrict user permissions ▪ Restrict log-in times ▪ Disable guest account ▪ Use failed attempts lockout ▪ Use timeout/screen lock ▪ Apply account expiration dates ▪ Change default administrator's user account/password ▪ Disable AutoRun ▪ Disable unused services	5
2.8 Given a scenario, apply common methods for securing mobile devices. ▪ Hardening techniques ▪ Device encryption ▪ Screen locks ▪ Facial recognition	6

I Introduction

Objective	Chapter(s)
▪ PIN codes ▪ Fingerprint ▪ Pattern ▪ Swipe ▪ Configuration profiles ▪ Patch management ▪ OS updates ▪ Application updates ▪ Endpoint security software ▪ Antivirus ▪ Antimalware ▪ Content filtering ▪ Locator applications ▪ Remote wipes ▪ Remote backup applications ▪ Failed log-in attempts restrictions ▪ Policies and procedures ▪ MDM ▪ BYOD vs. corporate-owned devices ▪ Profile security requirements	
2.9 Compare and contrast common data destruction and disposal methods.	5
▪ Physical destruction of hard drives ▪ Drilling ▪ Shredding ▪ Degaussing ▪ Incineration ▪ Recycling or repurposing best practices ▪ Erasing/wiping ▪ Low-level formatting ▪ Standard formatting ▪ Outsourcing concepts ▪ Third-party vendor ▪ Certification of destruction/recycling ▪ Regulatory and environmental requirements	

Objective	Chapter(s)
2.10 Given a scenario, apply security settings on SOHO wireless and wired networks.	6
▪ Router settings ▪ Change default passwords ▪ IP filtering ▪ Firmware updates ▪ Content filtering ▪ Physical placement/secure locations ▪ Universal Plug and Play (UPnP) ▪ Screened subnet ▪ Configure secure management access ▪ Wireless specific ▪ Changing the service set identifier (SSID) ▪ Disabling SSID broadcast ▪ Encryption settings ▪ Configuring guest access ▪ Firewall settings ▪ Disabling unused ports ▪ Port forwarding/mapping	
2.11 Given a scenario, configure relevant security settings in a browser.	6
▪ Browser download/installation ▪ Trusted sources ▪ Hashing ▪ Untrusted sources ▪ Browser patching ▪ Extensions and plug-ins ▪ Trusted sources ▪ Untrusted sources	

Objective	Chapter(s)
▪ Password managers ▪ Secure connections/sites–valid certificates ▪ Settings ▪ Pop-up blocker ▪ Clearing browsing data ▪ Clearing cache ▪ Private-browsing mode ▪ Sign-in/browser data synchronization ▪ Ad blockers ▪ Proxy ▪ Secure DNS ▪ Browser feature management ▪ Enable/disable ▪ Plug-ins ▪ Extensions ▪ Features	
3.0 Software Troubleshooting	
3.1 Given a scenario, troubleshoot common Windows OS issues.	7
▪ Blue screen of death (BSOD) ▪ Degraded performance ▪ Boot issues ▪ Frequent shutdowns ▪ Services not starting ▪ Applications crashing ▪ Low memory warnings ▪ USB controller resource warnings ▪ System instability ▪ No OS found ▪ Slow profile load ▪ Time drift	

Objective	Chapter(s)
3.2 Given a scenario, troubleshoot common mobile OS and application issues.	7
▪ Application fails to launch ▪ Application fails to close/crashes ▪ Application fails to update ▪ Application fails to install ▪ Slow to respond ▪ OS fails to update ▪ Battery life issues ▪ Random reboots ▪ Connectivity issues ▪ Bluetooth ▪ Wi-Fi ▪ Near-field communication (NFC) ▪ Screen does not autorotate	
3.3 Given a scenario, troubleshoot common mobile OS and application security issues.	7
▪ Security concerns ▪ Application source/unofficial application stores ▪ Developer mode ▪ Root access/jailbreak ▪ Unauthorized/malicious application ▪ Application spoofing ▪ Common symptoms ▪ High network traffic ▪ Degraded response time ▪ Data-usage limit notification ▪ Limited Internet connectivity ▪ No Internet connectivity ▪ High number of ads ▪ Fake security warnings ▪ Unexpected application behavior ▪ Leaked personal files/data	

Objective	Chapter(s)
3.4 Given a scenario, troubleshoot common personal computer (PC) security issues. ▪ Common symptoms ▪ Unable to access the network ▪ Desktop alerts ▪ False alerts regarding antivirus protection ▪ Altered system or personal files ▪ Missing/renamed files ▪ Inability to access files ▪ Unwanted notifications within the OS ▪ OS updates failures ▪ Browser-related symptoms ▪ Random/frequent pop-ups ▪ Certificate warnings ▪ Redirection ▪ Degraded browser performance	7

4.0 Operational Procedures

4.1 Given a scenario, implement best practices associated with documentation and support systems information management. 10

- Ticketing systems
 - User information
 - Device information
 - Description of issues
 - Categories
 - Severity
 - Escalation levels
 - Clear, concise written communication
 - Issue description
 - Progress notes
 - Issue resolution

Objective	Chapter(s)
▪ Asset management ▪ Inventory lists ▪ Configuration management database (CMDB) ▪ Asset tags and IDs ▪ Procurement life cycle ▪ Warranty and licensing ▪ Assigned users ▪ Types of documents ▪ Incident reports ▪ Standard operating procedures (SOPs) ▪ Software package custom installation procedure ▪ New user/onboarding setup checklist ▪ User off-boarding checklist ▪ Service-level agreements (SLAs) ▪ Internal ▪ External/third-party - Knowledge base/articles	
4.2 Given a scenario, apply change management procedures.	10
▪ Documented business processes ▪ Rollback plan ▪ Backup plan ▪ Sandbox testing ▪ Responsible staff members ▪ Change management ▪ Request forms ▪ Purpose of the change ▪ Scope of the change ▪ Change type ▪ Standard change	

Objective	Chapter(s)
▪ Normal change ▪ Emergency change ▪ Date and time of change ▪ Change freeze ▪ Maintenance windows ▪ Affected systems/impact ▪ Risk analysis ▪ Risk level ▪ Change board approvals ▪ Implementation ▪ Peer review ▪ End-user acceptance	
4.3 Given a scenario, implement workstation backup and recovery methods.	10
▪ Backup ▪ Full ▪ Incremental ▪ Differential ▪ Synthetic full ▪ Recovery ▪ In-place/overwrite ▪ Alternative location ▪ Backup testing ▪ Frequency ▪ Backup rotation schemes ▪ On-site vs. off-site ▪ Grandfather-father-son (GFS) ▪ 3-2-1 backup rule	

Objective	Chapter(s)
4.4 Given a scenario, use common safety procedures.	9
▪ Electrostatic discharge (ESD) straps ▪ ESD mats ▪ Electrical safety ▪ Equipment grounding ▪ Proper component handling and storage ▪ Cable management ▪ Antistatic bags ▪ Compliance with government regulations ▪ Personal safety ▪ Disconnect power before repairing PC ▪ Lifting techniques ▪ Fire safety ▪ Safety goggles ▪ Air filter mask	
4.5 Summarize environmental impacts and local environment controls.	9
▪ Material safety data sheet (MSDS) documentation for handling and disposal ▪ Proper battery disposal ▪ Proper toner disposal ▪ Proper disposal of other devices and assets ▪ Temperature, humidity-level awareness, and proper ventilation ▪ Location/equipment placement ▪ Dust cleanup ▪ Compressed air/vacuums ▪ Power surges, brownouts, and blackouts ▪ Uninterruptible power supply (UPS) ▪ Surge suppressor	

Objective	Chapter(s)
4.6 Explain the importance of prohibited content/activity and privacy, licensing, and policy concepts. ▪ Incident response ▪ Chain of custody ▪ Informing management/law enforcement as necessary ▪ Copy of drive (data integrity and preservation) ▪ Incident documentation ▪ Order of volatility ▪ Licensing/digital rights management (DRM)/end-user license agreement (EULA) ▪ Valid licenses ▪ Perpetual license agreement ▪ Personal-use license vs. corporate-use license ▪ Open-source license ▪ Non-disclosure agreement (NDA)/mutual non-disclosure agreement (MNDA) ▪ Regulated data ▪ Credit card payment information ▪ Personal government-issued information ▪ PII ▪ Healthcare data ▪ Data retention requirements ▪ Acceptable use policy (AUP) ▪ Regulatory and business compliance requirements ▪ Splash screens	9
4.7 Given a scenario, use proper communication techniques and professionalism. ▪ Present a professional appearance and wear appropriate attire. ▪ Match the required attire of the given environment. ▪ Formal ▪ Business casual	10

Objective	Chapter(s)
▪ Use proper language and avoid jargon, acronyms, and slang, when applicable. ▪ Maintain a positive attitude/project confidence. ▪ Actively listen and avoid interrupting the customer. ▪ Be culturally sensitive. ▪ Use appropriate professional titles and designations, when applicable. ▪ Be on time (if late, contact the customer). ▪ Avoid distractions. ▪ Personal calls ▪ Texting/social media sites ▪ Personal interruptions ▪ Appropriately deal with difficult customers or situations. ▪ Do not argue with customer and/or be defensive. ▪ Avoid dismissing customer issues. ▪ Avoid being judgmental. ▪ Clarify customer statements (i.e., ask open-ended questions to narrow the scope of the issue, restate the issue, or question to verify understanding). ▪ Use discretion and professionalism when discussing experiences/encounters. ▪ Set and meet expectations/timeline and communicate status with the customer. ▪ Offer repair/replacement options, as needed. ▪ Provide proper documentation on the services provided. ▪ Follow up with customer/user at a later date to verify satisfaction. ▪ Appropriately handle customers' confidential and private materials. ▪ Located on a computer, desktop, printer, etc.	

Objective	Chapter(s)
4.8 Explain the basics of scripting.	8
▪ Script file types ▪ .bat ▪ .ps1 ▪ .vbs ▪ .sh ▪ .js ▪ .py ▪ Use cases for scripting ▪ Basic automation ▪ Restarting machines ▪ Remapping network drives ▪ Installation of applications ▪ Automated backups ▪ Gathering of information/data ▪ Initiating updates ▪ Other considerations when using scripts ▪ Unintentionally introducing malware ▪ Inadvertently changing system settings ▪ Browser or system crashes due to mishandling of resources	
4.9 Given a scenario, use remote access technologies.	8
▪ Methods/tools ▪ RDP ▪ VPN ▪ Virtual network computer (VNC) ▪ Secure Shell (SSH) ▪ Remote monitoring and management (RMM) ▪ Simple Protocol for Independent Computing Environments (SPICE) ▪ Windows Remote Management (WinRM) ▪ Third-party tools ▪ Screen-sharing software	

Objective	Chapter(s)
▪ Videoconferencing software ▪ File transfer software ▪ Desktop management software ▪ Security considerations of each access method	
4.10 Explain basic concepts related to artificial intelligence (AI).	8
▪ Application integration ▪ Policy ▪ Appropriate use ▪ Plagiarism ▪ Limitations ▪ Bias ▪ Hallucinations ▪ Accuracy ▪ Private vs. public ▪ Data security ▪ Data source ▪ Data privacy	



Exam objectives are subject to change at any time without prior notice at CompTIA's sole discretion. Please visit CompTIA's website (www.comptia.org) for the most current listing of exam objectives.

How to Contact the Publisher

If you believe you have found a mistake in this book, please bring it to our attention. At John Wiley & Sons, we understand how important it is to provide our customers with accurate content, but even with our best efforts an error may occur.

In order to submit your possible errata, please email it to our Customer Service Team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Assessment Test for Exam 220-1202

1. Which will accomplish a specific task for the user?
 - A. Driver
 - B. Application
 - C. Operating system
 - D. Filesystem
2. Which operating system is a workstation that operates solely in the cloud?
 - A. Linux
 - B. Android
 - C. Windows 11
 - D. ChromeOS
3. Which term defines the precompiled instructions that are programmed and allows an operating system and its applications to operate?
 - A. Shell
 - B. Source
 - C. GUI
 - D. Multithreading
4. Which term best describes the Android operating system?
 - A. Server
 - B. Workstation
 - C. Mobile
 - D. Cloud-based
5. Which Windows Update branch will allow you to install preview releases of updates for Windows 11?
 - A. Semi-Annual Channel
 - B. Semi-Annual Channel (Targeted)
 - C. Long-Term Servicing Channel
 - D. Insider Program
6. Which tool will allow you to ready the operating system for imaging?
 - A. Microsoft Deployment Toolkit
 - B. Windows Assessment and Deployment Kit
 - C. sysprep.exe
 - D. Windows Imaging

7. Which element of the boot process holds the information that instructs the operating system to load from a specific partition?
 - A. `winload.efi`
 - B. `bootmgr`
 - C. `winresume.exe`
 - D. BCD
8. You need to configure a static IP address for Windows 11. Which of the following will allow you to perform this task?
 - A. The Network & Sharing Center
 - B. Windows Defender Firewall
 - C. The Network & Internet Settings screen
 - D. The VPN settings screen
9. Which filesystem performs on-the-fly defragmentation?
 - A. FAT
 - B. NTFS
 - C. `ext4`
 - D. FAT32
10. Which place can you find updates for Apple operating systems?
 - A. iTunes
 - B. App Store
 - C. Keychain
 - D. Mission Control
11. Which macOS feature is similar to Windows File Explorer?
 - A. Keychain
 - B. iCloud
 - C. Spotlight
 - D. Finder
12. You want to find all files in the `/home` directory and subdirectories that contain the word `password`. Which command should you use?
 - A. `find /home -name "password"`
 - B. `grep -r "password" /home`
 - C. `ls /home | grep "password"`
 - D. `grep "password" /home`

- 13.** Your system log files report an ongoing attempt to gain access to a single account. This attempt has been unsuccessful to this point. What type of attack are you most likely experiencing?
- A.** Password-guessing attack
 - B.** Rootkit attack
 - C.** Worm attack
 - D.** TCP/IP hijacking
- 14.** One of the vice presidents of the company calls a meeting with the information technology department after a recent trip to competitors' sites. She reports that many of the companies she visited granted access to their buildings only after fingerprint scans, and she wants similar technology employed at this company. Of the following, which technology relies on a physical attribute of the user for authentication?
- A.** Smartcard
 - B.** Biometrics
 - C.** Geo-fencing
 - D.** Tokens
- 15.** What type of malware is able to conceal itself from many parts of the operating system and will obtain/retain elevated privileges?
- A.** Worm
 - B.** Trojan
 - C.** Rootkit
 - D.** Botnet
- 16.** Which component of physical security addresses outer-level access control?
- A.** Fences
 - B.** Access control vestibule
 - C.** Multifactor authentication
 - D.** Strong passwords
- 17.** Which type of device can detect weapons on a person entering a facility?
- A.** Biometrics
 - B.** Magnetometer
 - C.** Motion sensor
 - D.** Badge reader
- 18.** A user wants to ensure that all contacts from their mobile device are copied onto their computer before migrating to a new device. Which of the following processes can accomplish this?
- A.** Mirroring
 - B.** Synchronization
 - C.** Calling each contact
 - D.** Attaching the contacts to an email

- 19.** Which feature of Windows Defender Firewall allows you to create custom rules for applications?
- A.** Real-time protection
 - B.** Cloud-delivered protection
 - C.** Advanced Security MMC
 - D.** Tamper protection
- 20.** A technician is configuring a new Windows computer for a home office. Which of the following steps should the technician take to secure the workstation?
- A.** Rename default accounts.
 - B.** Configure Single Sign-On.
 - C.** Disable Windows Update.
 - D.** Disable Action Center pop-ups.
- 21.** You need to protect files on the desktop operating system with encryption, but find out you do not have a TPM. What can you use to protect the files?
- A.** BitLocker
 - B.** Encrypted File System
 - C.** BitLocker To Go
 - D.** Full drive encryption
- 22.** Which is a benefit of setup of a device with a Microsoft account?
- A.** Access to the online store
 - B.** Automatic synchronization of OneDrive
 - C.** Active Directory authentication
 - D.** Offline authentication
- 23.** Which critical system generally requires third-party tools to troubleshoot performance problems?
- A.** RAM
 - B.** CPU
 - C.** Graphics
 - D.** Network
- 24.** What is the last step that should be performed when removing malware?
- A.** Investigate and verify symptoms.
 - B.** Enable System Protection.
 - C.** Educate the end user.
 - D.** Schedule scans and run updates.

- 25.** Your mobile device is suffering from an intermittent wireless connection. What is recommended to alleviate signal drops?
- A.** Shortening the SSID
 - B.** Using the 5 GHz band
 - C.** Reducing power on the WAP
 - D.** Using the 2.4 GHz band
- 26.** You are receiving USB controller resource warning messages. What should you try to resolve the issue?
- A.** Move USB devices around on the USB ports.
 - B.** Manually allot more endpoints.
 - C.** Upgrade drivers for the USB devices.
 - D.** Manually increase the output amperage.
- 27.** Why is time drift a big problem on virtual machines?
- A.** Lack of configuration for the NTP server
 - B.** Availability of the NTP server
 - C.** Shared physical RTC
 - D.** Emulated RTC
- 28.** Which tool can be used to restart a failed service?
- A.** msconfig.exe
 - B.** WinRE
 - C.** Computer Management MMC
 - D.** Resource Monitor
- 29.** Which is a safety risk from overheating a lithium-ion battery?
- A.** Reduced voltage
 - B.** Shock
 - C.** Explosion
 - D.** Shutdown
- 30.** A user is experiencing frequent shutdowns on their Windows computer. What should you check first to identify the cause?
- A.** Event Viewer
 - B.** Device Manager
 - C.** Task Manager
 - D.** Windows Update

31. Which network protocol and port does RDP operate on?
- A. TCP port 3389
 - B. TCP port 22
 - C. TCP port 5900
 - D. TCP port 443
32. Which line would be used to comment Windows Batch script code?
- A. //comment
 - B. 'comment
 - C. REM comment
 - D. # comment
33. Which command will launch the Remote Desktop Connection utility?
- A. msra.exe
 - B. mstsc.exe
 - C. quickassist.exe
 - D. ssh.exe
34. A client computer connects to the main office and is configured with an IP address from the main office, on the client computer. What is being described?
- A. Site-to-site VPN
 - B. Remote Desktop connection
 - C. SSH connection
 - D. Host-to-site VPN
35. Zoom and Teams are classic examples of what?
- A. Screen-sharing software
 - B. Videoconferencing software
 - C. File transfer software
 - D. Desktop management software
36. Which of the following is a hallmark of human intelligence that AI does not possess?
- A. Pattern recognition
 - B. Emotional experience
 - C. Data processing
 - D. Task automation
37. Which of the following are good measures to take to help prevent ESD? (Choose two.)
- A. Decrease the humidity.
 - B. Tie long hair back.
 - C. Take your shoes off.
 - D. Perform self-grounding.

38. Which of the following screwdrivers has a splined head?
- A. Flat-tipped
 - B. Phillips
 - C. Axial
 - D. Torx
39. What type of software licensing agreement usually provides free trial software, with the expectation that you will pay for it if you decide to keep it?
- A. Freeware
 - B. Shareware
 - C. Open source
 - D. Single user
40. How do antistatic mats and wrist straps drain the potential ESD?
- A. The use of magnets
 - B. Through a resistor
 - C. High-voltage probe
 - D. Direct ground
41. When cleaning dust out of computer equipment, you should always use which of the following?
- A. Multimeter
 - B. Flashlight
 - C. Air filter mask
 - D. Mirror
42. Which is the most likely source of interference to wireless communications?
- A. Magnets
 - B. ESD
 - C. Surges
 - D. Microwave ovens
43. While working on a user's system, you discover a sticky note attached to the bottom of the keyboard that has their username and password written on it. The user is not around, and you need to verify that the network connection is working. What should you do?
- A. Log in, verify access, and log out.
 - B. Log in and stay logged in when you are finished.
 - C. Text the user.
 - D. Log in and change the user's password.
44. While installing a new network card, you accidentally broke a component off of the card. What should you do?
- A. Explain that the card was broken out of the box.
 - B. Install the card anyway.

- C.** Inform the customer of the situation.
 - D.** Leave everything as is until you can locate a new card.
- 45.** What should you do when you cannot meet the timeline you proposed for resolving an issue?
 - A.** Avoid informing the customer.
 - B.** Communicate the delay and provide a new estimate.
 - C.** Blame external factors.
 - D.** Close the ticket prematurely.
- 46.** Which ticket entry method is the easiest for end users to submit their own tickets?
 - A.** Email entry
 - B.** Portal entry
 - C.** Manual entry
 - D.** Application entry
- 47.** Which method is used to test a change in an isolated environment?
 - A.** Primary plan
 - B.** Backout plan
 - C.** Sandbox testing
 - D.** Technical evaluation
- 48.** You need to change the computer's name and join it to a domain, which tool will allow you to perform these actions?
 - A.** Device Manager
 - B.** User Accounts
 - C.** System Properties
 - D.** Credential Manager
- 49.** Which tab in the Task Manager will allow you to see processes separated by each user on the system?
 - A.** Processes
 - B.** Performance
 - C.** App History
 - D.** Users
- 50.** Which RAID system requires three or more disks to provide fault tolerance?
 - A.** Mirroring
 - B.** Striping
 - C.** RAID-1
 - D.** Striping with Parity

Answers to Assessment Test 220-1202

1. B. An application is written to perform a specific task for the user of the operating system, such as word processing or data analysis. Drivers allow the operating system to communicate with the hardware. The operating system and the filesystem assist in running applications. See Volume 2 Chapter 1 for more information.
2. D. The Google ChromeOS is a workstation that operates solely in the cloud, because all the data is stored in the cloud. Cloud-based operating systems require an Internet connection to access information. Android and Windows 11 use cloud storage, but they are not considered cloud-based operating systems. Linux is strictly a workstation operating system. See Volume 2 Chapter 1 for more information.
3. B. The source, also known as the source code, is the code to be compiled, and this allows the operating system and applications to execute. The shell is a component of the operating system that allows execution of applications. The GUI is a graphical extension of the shell to allow the execution of applications. Multithreading is a term used to describe how an application can create multiple simultaneous requests to the processor. See Volume 2 Chapter 1 for more information.
4. C. The Android operating system is a mobile operating system often used with tablets and smartphones. A server is an operating system that is optimized to serve information. A workstation is an operating system that is optimized to execute and display the information retrieved from a server operating system. A cloud-based operating system is an operating system that requires network connectivity to access applications. See Volume 2 Chapter 1 for more information.
5. D. The Insider Program allows for the installation of brand-new features before they are publicly released. The Semi-Annual Channel is normally delayed by three to four months. The Semi-Annual Channel (Targeted) branch will install updates as they are released to the general public. The Long-Term Servicing Channel never installs new features during the life of the version of Windows. See Volume 2 Chapter 3 for more information.
6. C. The `sysprep.exe` utility will allow you to ready the operating system for imaging, by resetting specific information, such as the computer name. The Microsoft Deployment Toolkit can assist in creating the steps, but it calls on the `sysprep` tool. The Windows Assessment and Deployment Kit allows you to customize the Windows operating system for imaging, but it does not ready the operating system for imaging. Windows Imaging (WIM) is a file format to contain the image. See Volume 2 Chapter 3 for more information.
7. D. The Boot Configuration Data (BCD) holds the information that instructs the Windows Boot Manager (`bootmgr`) to load the operating system from a specific partition. `winload.efi` loads the operating system kernel. `bootmgr` is the initial bootstrap program that reads the BCD. `winresume.exe` is used when resuming a previous session that has been suspended. See Volume 2 Chapter 3 for more information.

8. C. By clicking the Start menu, then Settings, and choosing Network & Internet you can configure the properties of the network adapter to setup a static IP address. The Network & Sharing Center is used to view the network connection firewall profile and connection status. The Windows Defender Firewall is used to protect the operating system from malicious network connectivity. The VPN settings screen allows you to configure a virtual private network (VPN) connection. See Volume 2 Chapter 3 for more information.
9. C. The ext4 filesystem does not suffer from fragmentation, because it performs on-the-fly defragmentation. FAT is a 16-bit filesystem that suffers from fragmentation. NTFS is a journaled filesystem that suffers from fragmentation. FAT32 is a 32-bit filesystem that suffers from fragmentation. See Volume 2 Chapter 4 for more information.
10. B. The App Store is where you can find updates for the Apple operating system. iTunes is used to purchase and download music, videos and other content. The Keychain is used to store credentials on behalf of the user. Mission Control is used to view all of the currently running applications. See Volume 2 Chapter 4 for more information.
11. D. The macOS Finder is the equivalent to the Windows File Explorer. The Keychain is a password management system for the operating system. iCloud is a cloud-based storage and backup service. Spotlight helps you find applications, documents and other files. See Volume 2 Chapter 4 for more information.
12. B. The `grep -r` command searches recursively for the word `password` in all files within the `/home` directory. The `find /home -name "password"` command will not recursively search the `/home` directory. The `ls /home | grep "password"` command will only return files with `password` in the name. The `grep "password" /home` command will not recursively search directories. See Volume 2 Chapter 4 for more information.
13. A. A password-guessing attack occurs when a user account is repeatedly attacked using a variety of passwords. A rootkit attack would not be immediately seen and not shown as an ongoing effort. A worm attack would not be visible as an ongoing attempt to gain access. TCP/IP hijacking is a form of an on-path (formerly man-in-the-middle) attack. See Volume 2 Chapter 5 for more information.
14. B. Biometrics relies on a physical characteristic of the user to verify identity. Biometric devices typically use either a hand pattern or a retinal scan to accomplish this. Smart-cards contain a private certificate key and are protected with a passphrase. Geo-fencing uses GPS coordinates to require that authentication only happens within those parameters. Tokens are rotating numerical keys that you must physically have with you. See Volume 2 Chapter 5 for more information.
15. C. Rootkits are software programs that have the ability to hide certain things from the operating system. A worm is malware that replicates itself and infects other system. Trojans are programs that enter a system or network under the guise of another program. A botnet is a group of infected computers that can be remotely controlled via a command and control server. See Volume 2 Chapter 5 for more information.
16. A. Fences are intended to delay or deter entrance into a facility's outer-level of access. Access control vestibules are used for mid-layer access control to prevent tailgating. Multifactor authentication (MFA) is used for mid- and inner-layer access control. Strong passwords are used for mid- and inner-layer access control. See Volume 2 Chapter 5 for more information.

17. B. A magnetometer, also known as a metal detector, can detect weapons on a person entering a facility. Biometrics are used to authenticate someone based on their physical traits. A motion sensor is used to detect motion. A badge reader is used to scan security ID badges. See Volume 2 Chapter 5 for more information.
18. B. Synchronization can copy all contacts, programs, email messages, pictures, music, and videos between a mobile device and a computer. Mirroring is a term reserved for when devices replicate themselves to similar devices. Calling the contacts won't achieve the desired result, and emailing contacts is inefficient and does not immediately result in placing the contacts on the computer. See Volume 2 Chapter 6 for more information.
19. C. The Advanced Security MMC allows you to create custom rules for applications in Windows Defender Firewall. Real-time protection provides real-time antivirus and antimalware protection. Cloud-delivered protection allows for the automatic submission of malware samples. Tamper protection notifies you if Windows Defender Antivirus is turned off or tampered with. See Volume 2 Chapter 6 for more information.
20. A. Renaming the default accounts on the new Windows computer is the easiest way to secure the operating system. Configuring Single Sign-On should only be performed if a resource is required that only uses SSO as authentication. Disabling Windows Update will do the opposite of securing the operating system. Disabling Action Center pop-ups will also do the opposite of securing the operating system. See Volume 2 Chapter 6 for more information.
21. B. Encrypted File System (EFS) will encrypt and protect files on an NTFS filesystem without the need of the TPM. BitLocker requires a TPM and therefore will not work. BitLocker To Go is used with removable media. Full drive encryption is used with mobile devices. See Volume 2 Chapter 6 for more information.
22. B. When you set up a device with a Microsoft account, you automatically synchronize files in OneDrive, along with other known folders. Access to the online store can be achieved by launching the Microsoft Store and logging in; it does not require the user to log into the operating system with the Microsoft account. Active Directory authentication is applicable if the workstation is joined to an Active Directory domain. Offline authentication is the default method if you are not signing in with a Microsoft account. See Volume 2 Chapter 6 for more information.
23. C. Graphics cards usually require third-party tools to diagnosis performance problems. RAM problems can be diagnosed inside the operating system from Task Manager. CPU problems can be diagnosed inside the operating system from Task Manager. Network problems can be diagnosed inside the operating system from Task Manager. See Volume 2 Chapter 7 for more information.
24. C. The last step in the malware removal process is the education of the end user to prevent future occurrences of infection. Investigate and verify malware symptoms is the first step in the malware removal process. Enabling System protection is one of the last steps, but it is not the last step. Scheduling scans and running updates should be performed after you have remediated the malware. See Volume 2 Chapter 7 for more information.
25. B. By using the 5 GHz wireless band you can limit the amount of interference from external devices such as microwave ovens and Bluetooth devices. Shortening the SSID will have no effect on intermittent signal drops. Reducing power on the WAP will have no effect on intermittent signal drops. Using the 2.4 GHz band will probably make signal drops worse. See Volume 2 Chapter 7 for more information.


- 26.** A. Moving USB devices around on the USB ports is the first step in remediating USB issues related to resources. USB 2.0 devices should be consolidated to USB 2.0 ports and the same should be done for USB 3.0 devices. You cannot manually allot endpoints for USB devices. Upgrading the drivers for the USB devices will not solve power issues or endpoint issues; drivers for the USB controller should be updated. You cannot manually increase the output amperage for USB devices. See Volume 2 Chapter 7 for more information.
- 27.** D. Normally a dedicated chip called the real time clock (RTC) keeps time for physical machines. However, with virtual machines their RTC is emulated and susceptible to time drift. By default, Windows 10/11 are configured for `time.windows.com` as an NTP. The availability of the NTP server of `time.windows.com` is not any different than physical machines. When deploying virtual machines, the physical RTC is not shared; each VM gets an emulated RTC. See Volume 2 Chapter 7 for more information.
- 28.** C. The Computer Management MMC contains the Services add-in that can be used to restart and configure failed services. Although services can be viewed with the `msconfig .exe` tool, they cannot be restarted with the tool, only enabled and disabled on startup. The Windows Recovery Environment (WinRE) is used to troubleshoot and repair problems offline. The Resource Monitor cannot be used to restart services. See Volume 2 Chapter 7 for more information.
- 29.** C. The risk of thermal runaway is attributed from overheating a lithium-ion battery. When thermal runaway happens, you run the risk of explosion of the battery. Overheating a li-ion battery can cause reduced voltage, but that is not a safety risk. Overheating a li-ion battery will not cause a shock. Overheating of the li-ion battery will cause a shutdown, but that is not a safety risk. See Volume 2 Chapter 7 for more information.
- 30.** A. Event Viewer logs can help identify the cause of frequent shutdowns by showing entries related to Kernel-Boot or Kernel-General. Device Manager manages hardware devices. Task Manager monitors tasks and performance. Windows Update manages updates for the operating system. See Volume 2 Chapter 7 for more information.
- 31.** A. The Remote Desktop Protocol operates on TCP port 3389. The SSH protocol operates on TCP port 22. The VNC protocol operates on TCP port 5900, and HTTPS operates on TCP port 443. See Volume 2 Chapter 8 for more information.
- 32.** C. The line of `REM comment` is used to comment Windows Batch script code. The line of `//comment` is used to comment JavaScript code. The line of `'comment` is used to comment VBScript code. The line of `# comment` is used to comment Bash script code and PowerShell code. See Volume 2 Chapter 8 for more information.
- 33.** B. The command `mstsc.exe` will launch the Windows Remote Desktop Connection utility. From this utility you can remotely connect to a server or other workstation. The command `msra.exe` will launch the Microsoft Remote Assistance utility to allow a trusted helper to remote in to help. The command `quickassist.exe` will launch the Quick Assist remote assistance utility to allow an assistant to remote in to help. The command `ssh.exe` will launch the Secure Shell client that allows you to connect to a Linux/UNIX server or networking equipment. See Volume 2 Chapter 8 for more information.

- 34.** D. The scenario being described is a host-to-site connection. A site-to-site VPN connection will be built between two sites, such as the main site and a branch site. A Remote Desktop connection or SSH connection will not configure a site IP address on the client computer. See Volume 2 Chapter 8 for more information.
- 35.** B. Zoom and Teams are classic examples of videoconferencing software, because they allow multiple attendees and have a form of attendee management. Microsoft Quick Assist is an example of screen-sharing software. Dropbox is an example of file transfer software. Splashtop and TeamViewer are examples of desktop management software. See Volume 2 Chapter 8 for more information.
- 36.** B. Emotional experience is a hallmark of human intelligence that AI does not possess. AI can recognize patterns. AI is excellent at processing large volumes of data. AI can automate tasks. See Volume 2 Chapter 8 for more information.
- 37.** B, D. Long hair or neckties can very easily hold a static charge and damage computer parts. In addition, if you don't have an ESD strap or mat, you should perform self-grounding before you touch sensitive computer equipment. Lower humidity actually increases the risk of ESD. Taking your shoes off will not limit ESD. See Volume 2 Chapter 9 for more information.
- 38.** D. A Torx screwdriver has a splined head used for greater gripping of the screw. A flat-tipped screwdriver as the name suggests has a flat tip. A Phillips screwdriver has a cross used for better gripping over a normal flat-tip screwdriver. An axial screwdriver is not a type of screwdriver. See Volume 2 Chapter 9 for more information.
- 39.** B. Shareware often provides a free trial of a program, with the expectation that you will pay for it if you keep it. Freeware and open source applications are generally totally free, although the developers may ask for donations. Single user software is licensed for one user. See Volume 2 Chapter 9 for more information.
- 40.** B. Typically, a 1 mega-ohm resistor is used to slowly drain or discharge the potential electrostatic discharge (ESD). Magnets are used to hold screws onto screwdrivers. High-voltage probes are used to drain high voltage stored up in electrical circuits and not ESD. Direct grounds are never used because they have the potential for discharge and possible electrocution. See Volume 2 Chapter 9 for more information.
- 41.** C. An air filter mask will protect you from inhaling dust particles. A multimeter is used to measure voltage. Although useful, a flashlight and a mirror are not needed for cleaning dust out of computer equipment. See Volume 2 Chapter 9 for more information.
- 42.** D. Microwave ovens operate in the 2.45 GHz range and can directly affect wireless communications. Magnets can generate a magnetic field, but they will not disrupt wireless communications. Electrostatic discharge (ESD) can kill components, but a discard from ESD in the same space as wireless will not cause interruptions. Electrical surges can affect the equipment running wireless, but it will not affect wireless communications. See Volume 2 Chapter 9 for more information.
- 43.** C. You should text the user and let them know that you need to verify their network connection. You should not log in as another person, with their username and password. You should not log in and stay logged in when you are finished. You should definitely not log in and change the user's password. See Volume 2 Chapter 10 for more information.

- 44.** C. You should be honest and inform the customer of the situation. Explaining the card was broken out of the box is untruthful and deceptive. Installing the card anyway will cause new problems. Leaving everything as is until you can locate a new card, does not help the customer since you've created a new dilemma. See Volume 2 Chapter 10 for more information.
- 45.** B. Communicating the delay and providing a new estimate will help maintain transparency and manage customer expectations. Avoiding informing the customer will not meet the customer's expectations for communicating the timeline. Blaming external factors is deceitful and not professional. Closing the ticket prematurely is an unprofessional behavior because the customer will not have a status on the resolution. See Volume 2 Chapter 10 for more information.
- 46.** A. Email entry is the easiest for end users because it allows them to email the ticketing system and use a method of communication they are familiar with. Portal entry requires the end user to log into the portal, then enter fields in a form to submit a ticket. Manual ticket entry is not done by the end user; it is completed by help desk staff. Application entry is an automated ticket that submits when an application has an issue; this is independent from what the user needs solved. See Volume 2 Chapter 10 for more information.
- 47.** C. Sandbox testing is used to test a change before placing it into production. The process from the sandbox testing will become the primary plan used when the change in production is green-lighted. The backout, also known as the rollback plan is created in the event the primary plan fails in production. The technical evaluation is a technical evaluation of the change, along with a draft of the primary plan. See Volume 2 Chapter 10 for more information.
- 48.** C. The System Properties applet (SYSDM.CPL) allows you to change the computer name and join the system to a domain. The Device Manager is used to manage hardware resources. The User Accounts applet is used to manage user accounts. Credential Manager is used to manage stored credentials. See Volume 2 Chapter 2 for more information.
- 49.** D. The Users tab in Task Manager will allow you to see processes sorted by each user. The Processes tab will show you processes sorted by forefront and background applications. The Performance tab will show you the overall performance of the system with graphical charts. The App History tab will display historical information about processes, such as CPU time. See Volume 2 Chapter 2 for more information.
- 50.** D. Striping with parity is also known as RAID-5, it requires three or more disks and provides fault tolerance. Mirroring is also known as RAID-1 and only requires two disks. Striping is also known as RAID-0 and provides no fault tolerance. See Volume 2 Chapter 2 for more information.

CompTIA A+ Core 2 Exam (220-1202)

**VOLUME
2**





Chapter 1

Operating System Basics

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **1.0 Operating Systems**
- ✓ **1.1 Explain common operating system (OS) types and their purposes.**
 - Workstation systems (OSs)
 - Windows
 - Linux
 - macOS
 - Chrome OS
 - Mobile OSs
 - iPadOS
 - iOS
 - Android
 - Vendor life-cycle limitations
 - End-of-life (EOL)
 - Update limitations
 - Compatibility concerns between operating systems
- ✓ **1.2 Given a scenario, perform OS installations and upgrades in a diverse environment.**
 - Upgrade considerations
 - Backup files and user preferences
 - Application and driver support/backward compatibility
 - Hardware compatibility
 - Feature updates
 - Product life cycle



✓ **1.3 Compare and contrast basic features of Microsoft Windows editions.**

- Windows 10 editions
 - Home
 - Pro
 - Pro for Workstations
 - Enterprise
- Windows 11 editions
 - Home
 - Pro
 - Enterprise
- N versions
- Feature differences
 - Domain vs. workgroup
 - Desktop styles/user interface
 - Availability of Remote Desktop Protocol (RDP)
 - Random-access memory (RAM) support limitations
 - BitLocker
 - gpedit.msc
- Upgrade paths
 - In-place upgrade
 - Clean install
- Hardware requirements
 - Trusted Platform Module (TPM)
 - Unified Extensible Firmware Interface (UEFI)

✓ **1.10 Given a scenario, install applications according to requirements.**

- System requirements for applications
 - 32-bit vs. 64-bit dependent application requirements
 - Dedicated vs. integrated graphics card



- Video random-access memory (VRAM) requirements
 - RAM requirements
 - Central processing unit (CPU) requirements
 - External hardware tokens
 - Storage requirements
 - Application to OS compatibility
 - Distribution methods
 - Physical media vs. mountable ISO file
 - Downloadable package
 - Image deployment
 - Impact considerations for new applications
 - Device
 - Network
 - Operation
 - Business
- ✓ **1.11 Given a scenario, install and configure cloud-based productivity tools.**
- Email systems
 - Storage
 - Sync/folder settings
 - Collaboration tools
 - Spreadsheets
 - Videoconferencing
 - Presentation tools
 - Word processing tools
 - Instant messaging
 - Identity synchronization
 - Licensing assignment



The first half of this book focused mainly on the hardware and physical elements of the computing environment. You learned about physical components, or hardware, of personal computers, laptops, and mobile devices, as well as networking, printers, and troubleshooting procedures. That completes the coverage of the topics on the 220-1201 exam. This chapter marks a departure from that.

In this chapter—and several to come—the focus is on operating systems (OSs). To be specific, the majority of information will be on the Microsoft Windows operating systems, which you must know well for the 220-1202 certification exam. However, the 220-1202 exam also requires basic knowledge of the macOS and Linux operating systems.

Understanding Operating Systems

Computers are pretty much useless without software. A piece of hardware might just as well be used as a paperweight or doorstop if you don't have an easy way to interface with it. Software provides that way. While there are many types of software, or programs, the most important one you'll ever deal with is the operating system. Operating systems have many different, complex functions, but two of them jump out as being critical: interfacing with the hardware and providing a platform on which other applications can run.

Here are three major distinctions of software about which you should be aware:

Operating System The *operating system* (OS) provides a consistent environment for other software to execute commands. The OS provides users with an interface with the computer so that they can send commands (input) and receive feedback or results (output). To do this, the OS must communicate with the computer hardware to perform the following tasks, as illustrated in Figure 1.1:


- Disk storage and file management
- Device access
- Memory management
- Output formatting

Once the OS has organized these basic resources, users can give the computer instructions through input devices (such as a keyboard or a mouse). Some of these commands are built into the OS, whereas others are issued through the use of applications. The OS becomes the center through which the system hardware, other software, and the user communicate; the rest of the components of the system work together through the OS, which coordinates their communication.

Application Used to accomplish a particular task, an *application* is software that is written to supplement the commands available to a particular OS. Each application is specifically compiled (configured) for the OS on which it will run. For this reason, the application relies on the OS to do many of its basic tasks. Examples of applications include complex programs, such as Microsoft Word and Microsoft Edge, as well as simple programs, such as a command-line utility. Whether they are complex or simple, when accessing devices and memory, the programs can simply request that the OS do it for them. This arrangement saves substantially on programming overhead because much of the executable code is *shared*—that is, it is written into the operating system and can therefore be used by multiple applications running on that OS.

Drivers *Drivers* are extremely specific software written for the purpose of instructing a particular OS on how to access a piece of hardware. Each modem or printer has unique features and configuration settings, and the driver allows the OS to understand how the hardware works and what it is able to do.

FIGURE 1.1 The operating system interacts with resources.



NOTE

You can learn about drivers in detail in Volume 1 Chapter 2, “Expansion Cards, Storage Devices, and Power Supplies.”

In the following sections, we’ll look at some terms and concepts central to all operating systems. Then we’ll move into specific discussions of Microsoft Windows operating systems.

Are These the Only Versions of Windows?

In the workplace, it is likely that you will encounter some different versions of the Windows operating systems. Although older versions of Windows are no longer supported by Microsoft, these operating systems can still be found in certain instances. These older operating systems are typically used to support legacy systems, such as kiosks, medical equipment,

and many other applications. Windows 11 has rapidly become the mainstream operating system for both the workplace and home over the past several years, since its release in 2021. Microsoft has scheduled the end of support for Windows 10 in late 2025, as of the writing of this book.

Know that this version of the exam focuses only on the Windows 10/11 operating systems mentioned here, as well as macOS, Linux (both detailed in Volume 2 Chapter 4, “Working with macOS and Linux”), ChromeOS, iOS, and iPadOS. We highly recommend that you become familiar with all of the operating systems your job requires, but these are the ones that you will need to know for the A+ certification exam.



Exam Objective 1.1, “Explain common operating system (OS) types and their purposes,” requires you to know the workstation and mobile operating systems that you might encounter as a technician. Be sure to know the differences between each of the operating systems.

Operating System Terms and Concepts

Before we get too far into our discussion of PC workstation operating systems, it will be useful to define a few key terms. The following are some terms that you will come across as you study this chapter and work in the computer industry:

Version A *version* is a particular revision of a piece of software, normally described by a number that tells you how new the product is in relation to other versions of the product.

Source Code The *source code* is the actual code that defines how a piece of software works. Computer operating systems can be *open source*, meaning that the OS can be examined and modified by anyone, or they can be *closed source*, meaning that only an owner or developer can modify or examine the code. As an example, Microsoft and macOS are closed source operating systems and *Linux* is an open source operating system.



A word often used interchangeably with *closed source* is *proprietary*.

Shell A *shell* is a program that runs on top of the OS and allows the user to issue commands through a set of menus or another interface (which may or may not be graphical). Shells make an OS easier to use by changing the user interface. File Explorer is the name of the shell that has been used with the Windows operating system since Windows 95.

Graphical User Interface A *graphical user interface*, or *GUI*, is a method by which a person communicates with a computer using graphical images, icons, and methods other than text. GUIs allow a user to use a mouse, touchpad, or another mechanism (in addition to the keyboard) to interact with the computer to issue commands.

Network A *network* is any group of computers that have a communication link between them. Networks allow computers to share information and resources quickly and securely.

Cooperative Multitasking *Cooperative multitasking* is a multitasking method that depends on the application itself to be responsible for using the processor and then freeing it for access by other applications. This is the way very early versions of Windows managed multiple applications. If any application locked up while using the processor, the application was unable to free the processor to do other tasks and the entire system locked, usually forcing a reboot.

Preemptive Multitasking *Preemptive multitasking* is a multitasking method in which the OS allots each application a certain amount of processor time and then forcibly takes back control and gives another application or task access to the processor. This means that if an application crashes, the OS takes control of the processor away from the locked application and passes it on to the next application, which should be unaffected. Although unstable programs still lock, only the locked application will stall, not the entire system. Preemptive multitasking is what is used today in modern operating systems.

Multithreading *Multithreading* is the ability of a single application to have multiple requests in to the processor at one time. This results in faster application performance because it allows a program to do many things at once.

32-Bit A *32-bit operating system* is one that can not only run on a 32-bit processor but can utilize the capabilities of the processor fully. While this may sound simple, the truth of the matter is that it took many years after the 32-bit processor became available before operating systems (which were 16-bit at the time) were able to utilize their features. You cannot execute 64-bit software on 32-bit operating systems; the instructions sets are functionally different.

64-Bit A *64-bit operating system* is one that is written to utilize the instructions possible with 64-bit processors. Originally, these were more common with servers than desktops, but they have now become ubiquitous in the market with both Intel and AMD processors. As mentioned earlier, you cannot mix 64-bit software with 32-bit hardware (but you can run most 32-bit software on 64-bit hardware).



Exam Objective 1.10 specifically calls out 32-bit vs. 64-bit dependent application requirements.

x86 The term *x86* is commonly used to refer to operating systems intended to run on the Intel processor. Intel initially identified its 32-bit processors with numbers ending in 86 prior to switching to the Pentium line.

x64 The term *x64* is commonly used to denote operating systems that can run on 64-bit processors. This is also commonly referred to as *AMD64*, because AMD defined the 64-bit instruction set used today.

ARM The term *ARM* is commonly used to denote operating systems that can run on Advanced RISC Machine (ARM) processors. ARM operating systems use a reduced instruction set computing (RISC) instruction set commonly found on tablet and phone processors. ARMv8-A processors can support 64-bit operating systems, so you will find 32-bit and 64-bit operating systems.

Random Access Memory The term *random access memory (RAM)* is used to describe the amount of memory installed and accessible for an operating system. The operating system uses RAM to hold its operating system code, as well as the programs you open on the operating system. Each program you install will have a minimum RAM requirement and recommended RAM requirement. Most computer systems on the market today are sold with at least 4 gigabytes (GB) of RAM. PC manufacturers have identified that newer programs are requiring more and more RAM to operate properly, so upgrades of 8 GB of RAM are common and have become the new normal.

A 32-bit operating system has the limitation of addressing only 4 GB of RAM. Most new computer systems on the market today come with a preinstalled 64-bit operating system. The 4 GB limitation will present itself if you are re-installing a computer system that has more than 4 GB of RAM and use a 32-bit version of the operating system. You will find out quickly that you just downgraded the computer.



Operating system vendors implemented a Physical Address Extension (PAE) workaround for the 4 GB limit. For the purposes of the exam, however, you should be aware that 32-bit operating systems have a hard limit of addressing 4 GB of RAM.

Operating Systems

When we think of an operating system, Windows or macOS is probably the first that comes to mind. These operating systems comprise only one category of operating system. We use other operating systems, such as Android and Apple iOS, and don't even realize how much they have become part of our daily life. The CompTIA exam recognized these various operating systems and integrated them into the 220-1202 exam.

Operating System Categories

An operating system category defines the use and function of both the operating system and the hardware. All operating systems fit into one of four different broad categories: *server*, *workstation*, *mobile*, or *cloud-based* operating systems.

Server Server operating systems allow for the sharing of information and applications by workstation operating systems. The operating system is optimized for background processes used for access by clients, as opposed to the foreground processes, such as the GUI.

Workstation Workstation operating systems allow end users to access information and applications on server operating systems as well as independently run applications locally. The operating system is optimized for foreground processes, such as the GUI. Background sharing is usually limited to a specific number of users; workstation versions of Microsoft Windows are limited to 10 concurrent users.

Mobile Mobile operating systems are found on mobile devices such as phones and tablets. The mobile operating system is generally optimized for touch-based devices where one program is used at a time.

Cloud-Based Cloud-based operating systems are a new breed of operating system that has emerged from the mobile computing era. Cloud-based operating systems are also considered stateless operating systems, since personal data is not primarily stored on the device but in the cloud. Mobile operating systems can be considered cloud-based operating systems. So this category takes on a dual role of defining a category and describing where the data is stored.

Common Operating Systems

Now that you understand the various categories of operating systems you may encounter, we will explore the common operating systems you will find in your everyday work as an A+ technician.

Microsoft Windows The Microsoft Windows operating system is one of the most popular operating systems adopted by business and personal users. Windows 11, the current version, is considered a workstation operating system.

Apple Macintosh OS *macOS* is the current version of the operating system for Apple Macintosh laptop and computer hardware platforms. Ironically, macOS is the original operating system name from back in the 1980s. Between 1999 and 2015, the macOS was called Mac OS X. The macOS operating system is considered a workstation operating system.

Linux Linux was created as a free, open source operating system in 1991 by Linus Torvalds. Linux was created to mimic the design of the Minix (mini-UNIX) operating system, which was a minimal version of UNIX that was originally used for academic purposes. Today, the Linux kernel (core of the OS) remains a free, open source platform upon which many different Linux distributions have been created, such as Ubuntu, Red Hat, Fedora, and Debian, just to name a few. The Linux operating system is considered both a workstation operating system and a server operating system, depending on its installation and configuration.

Microsoft Windows Tablet In 2012, with the release of the tile-based Windows 8 operating system, Microsoft released the Surface tablet. Today the Surface tablet runs the Windows 11 operating system and uses an ARM-based processor. It competes on the market with other *Ultrabooks* (subnotebook computers), such as the iPad Pro. Newer versions of the Surface tablet even have a neural processing unit (NPU) for artificial

intelligence (AI) processing. The operating system itself is a full version of Windows 11, and it's considered a mobile operating system.

Android Google created the *Android* operating system as a tile-based mobile operating system. Android was originally released in 2008 with HTC phones; the release of Android tablets followed in 2011. Android-based phones and tablets have become quite popular in the mobile operating system market. The popularity of Android is attributed to the application ecosystem that the Google Play store creates. The Google Play store enables users to purchase, download, and install applications, books, music, and video for the Android operating system.

iOS Apple created the iOS operating system as a tile-based mobile operating system. Originally released on the first iPhone in 2007, the operating system was known as the iPhone Operating System (iOS). The release of the iPod Touch followed a few months later, and it also furnished the iPhone Operating System. In 2010, the first iPad was released, and it furnished the iPhone Operating System as well. Today, the operating system's name has been shortened to iOS; the current version, as of this writing, is iOS version 18. Apple iOS is a popular mobile operating system in the marketplace today. Its popularity is attributed to Apple's strong support for the accompanying hardware. The Apple App Store also creates a strong ecosystem for downloading, installing, and purchasing applications, books, music, and video for the iOS operating system.

iPadOS *iPadOS* version 13.1 was released in the fall of 2019 for the Apple iPad. iPadOS is a variant of the popular Apple iOS with better multitasking features and a better home screen. This allows the Apple iPad to behave more like a light-weight laptop than a tablet. The Apple iPad Pro can be expanded with a keyboard, and has built-in AI called Apple Intelligence. This makes it a real competitor in the Ultrabook market.

ChromeOS *ChromeOS*, formerly styled as Chrome OS, is an operating system that is designed around the Chrome web browser by Google. ChromeOS was released in 2011 on the hardware platform known as the Chromebook, which is considered one of the first Ultrabooks. ChromeOS is a true cloud-based operating system, since all its data is stored in the Google Drive cloud-based storage system. With the initial release of the Chromebook, workstation-type devices called Chromeboxes had also been released to market. Applications, which are really just web plug-ins, can be installed from the Chrome Web Store. The applications are written in HTML, HTML5, CSS, and JavaScript, which allow cross-platform support between ChromeOS and the Chrome web browser installed on a PC or Mac.

Stateless Operating Systems

The lines have blurred in the past few years with the rapid adoption of the cloud for both enterprise and consumer installations. The workstation-based operating systems, such as Windows 11 and macOS, have embraced the cloud and to some extent have become

stateless devices like mobile-based operating systems. Other devices like ChromeOS are strictly cloud-based devices and contain no local filesystem, yet act as workstations for cloud-based applications.

In the enterprise, an operating system can be set up with all the applications a user needs to do their job, along with access to all their files. The installation of applications and the configuration of the applications is managed by mobile device management (MDM) software, such as Intune. Then, if the device breaks, is stolen, or is upgraded to new hardware, a new device can be provisioned in the MDM software and everything gets set up just like their previous device.

Consumer installations have also benefited from the adoption of the cloud. When a consumer purchases a new device to replace their old device, all they have to do is sign into the cloud service and their files and applications will automatically be transferred over and installed. The adoption of the cloud and the new behavior now makes traditional operating systems act more like stateless mobile-based operating systems.



Exam Objective 1.2, “Given a scenario, perform OS installations and upgrades in a diverse environment.” Be sure to understand the operating system life cycle in the following section.

Operating System Life Cycle

An operating system life cycle begins when an operating system is introduced and ends when the operating system is no longer supported. As a computer technician, you should pay close attention to an operating system’s life cycle, because the *end-of-life (EOL)* date for an OS means that it will no longer receive updates.

When an operating system is considered end-of-life, newer features will not be added to the OS in the future. More importantly, security updates will no longer be offered, which will put your operating system and information at risk of compromise.

The network administrator should pay close attention to the dates an operating system is considered end-of-life. If the organization has a support contract for problems, the support of the operating system at the end-of-life could no longer be honored. Of course, the lack of security updates for a corporate network is more concerning, since the stability of an organization can be compromised.

It has been common practice for network administrators to skip versions of every other operating system. This is mainly attributed to the *long-term support* for the current operating system and the amount of work required to upgrade an entire organization’s operating systems. In addition to these factors, many hardware vendors support the current operating system plus the last release. Administrators can purchase a new laptop or PC and request the prior operating system. However, once the next operating system is released, the support for the old operating system will force an upgrade.



When Windows Vista was released, Windows XP was the dominant corporate operating system. Windows Vista was not highly adopted, but when Windows 7 was released, many administrators planned on the highly anticipated upgrade to Windows 7. We've seen the same trend with Windows 10; many administrators completely skipped rolling out Windows 8 entirely. Windows 11 is the current operating system and there is no way to tell what will happen in the future. Microsoft has changed its naming convention for Windows and released platform upgrades as feature releases. This has made it easier for administrators to adopt the new Windows version.

The following common terminology is used by operating system developers in relation to the life cycle of an OS:

Alpha An operating system in alpha is a first draft of an operating system. Operating systems that are available in alpha are considered previews. It is common to find these in open source communities, such as Linux operating systems, but it is uncommon to find a corporate-based operating system in alpha.

Beta An operating system in beta is a second draft of an operating system. This stage of an operating system means most of the bugs have been discovered and are avidly being patched. This stage is typically used to discover missing features and new bugs. Some operating system developers require an opt-in to run a beta operating system, and some require invites to a beta program that mandates participation.

Release Candidate An operating system that is deemed a release candidate is in the final stage before it is released to market. During the release candidate stage, administrators are urged to install and use the operating system as though it were the final release. The benefit is that the administrator can get a previewed look at the operating system while helping to polish the final product, also called the *release to market (RTM)*. An operating system in this stage is generally stable and has relatively few bugs.

Current A current operating system, as the term states, is the current release of an operating system. Many operating system developers are on a one- to two-year release cycle for the future version of the current operating system.

Service Pack Although the term *service pack* is generally associated with updates, it is also used to describe a milestone in the life cycle of an operating system. When a service pack is released for an operating system, it adds major features as well as patches for both security and functionality. Microsoft has replaced this term with the introduction of Windows 8; it is now just called an update or Windows Update. However, many other operating systems still use “service pack.”

Mainstream Support During the *mainstream support* of an operating system, all *hot-fixes*, *security updates*, *feature updates*, and general support for the operating system is supported. An operating system developer will usually have a predetermined mainstream support date when the operating system is released. When the mainstream support ends

for an operating system, it is either no longer supported or enters into an extended support period.

Extended Support The *extended support* period, sometimes called the *long-term support (LTS)* period, for an operating system is the final state of an OS's life cycle. For example, during this time, hotfixes and features will no longer be supported, but critical security updates will be supported with Microsoft operating systems. Additional contract purchases may be required if hotfixes are required during this period. An administrator should have all systems upgraded to the current operating system by the end of the extended support date.



Microsoft has announced end of support date of October 14, 2025 for Windows 10 Home and Pro. The end of support for Windows 10 (21H2) LTSC is January 12, 2027. You can read about other versions of Windows 10 at <https://learn.microsoft.com/en-us/windows/release-health/release-information>.



Exam Objective 1.3, "Compare and contrast basic features of Microsoft Windows editions," requires you to know the hardware requirements for Windows 10/11. Be sure to know the various hardware requirements.

Minimum System Requirements

In the upcoming chapters, we'll explore how to install and upgrade each of the operating systems that you need to know for the exam. However, the hardware requirements of the operating system that you are thinking of installing can prevent you from even considering these options. Before you can begin to install an OS, you must consider several items. You must perform the following tasks before you even start to think about undertaking the installation:

1. Determine hardware compatibility and *minimum requirements*.
2. Determine installation options.
3. Determine the installation method.
4. Prepare the computer for installation.

Let's begin our discussion by talking about hardware compatibility issues and the requirements for installing the various versions of Windows.

Determining Hardware Compatibility and Minimum Requirements

Before you can begin to install any version of an operating system, you must determine whether it supports the hardware that you will be using. That is, will the version of the operating system have problems running any of the device drivers for the hardware that you have?

To answer this question, operating system vendors have developed several versions of *hardware compatibility lists (HCLs)*. An HCL is a list of all the hardware that works with the operating system. Microsoft published HCLs for many prior versions of Windows. Since the release of Windows 10, the HCL has disappeared completely and it is now the responsibility of the vendor to certify the compatibility of the hardware with Microsoft. However, many other operating system vendors still have a recommended HCL.

The point is, before installing an operating system, you should check all your computer's components against an HCL or the manufacturer of the component to ensure that they are compatible with the version of operating system you plan to install. If a product is not on the list, that does not mean it will not work; it merely means it has not been tested. The list represents tested software and hardware that vendors have stated are compatible, but it is by no means all-inclusive.

In addition to general compatibility, it is important that your computer have enough “oomph” to run the version of an operating system that you plan to install. For that matter, it is important for your computer to have enough resources to run any software that you plan to use. Toward that end, Microsoft (as well as other software publishers) releases a list of both minimum and recommended hardware specifications that you should follow when installing Windows.

“Minimum specifications” are the absolute minimum requirements for hardware that your system should meet in order to install and run the OS version you have chosen. “Recommended hardware specifications” are what you should have in your system in order to realize usable performance. Always try to have the recommended hardware (or better) in your system. If you don't, you may have to upgrade your hardware before you upgrade your operating system if you're running anything beyond a minimal environment. Table 1.1 lists the minimum hardware specifications for Windows 10 and Windows 11. Note that in addition to these minimum requirements, the hardware chosen must be compatible with the selected version of Windows. Also be aware that additional hardware may be required if certain features are installed (for example, a fingerprint reader is required to use biometric logins). Windows 11 is only available in a 64-bit version and requires a Trusted Platform Module 2.0 and UEFI for added security of the operating system and its applications.

TABLE 1.1 Windows 10 and Windows 11 minimum system requirements

Operating System	Windows 10	Windows 10	Windows 11
Architecture	32-bit	64-bit	64-bit
Processor	1 GHz or faster processor or System on a Chip (SoC)	1 GHz or faster processor or System on a Chip (SoC)	1 GHz or faster processor with 2 or more cores or System on a Chip (SoC)
Memory	1 GB	2 GB	4 GB
Free hard disk space	16 GB	32 GB	64 GB

Operating System	Windows 10	Windows 10	Windows 11
Graphics card	Microsoft DirectX 9 or later graphics device with WDDM 1.0 driver	Microsoft DirectX 9 or later graphics device with WDDM 1.0 driver	Microsoft DirectX 12 or later graphics device with WDDM 2.0 driver
Display	800 × 600	800 × 600	High-definition 720p
Additional hardware	N/A	N/A	UEFI & TPM 2.0 required

If there is one thing to be learned from Table 1.1, it is that Microsoft is nothing if not optimistic. For your own sanity, though, we strongly suggest that you always take the minimum requirements with a grain of salt. They are, after all, *minimum* requirements. Even the recommended requirements should be considered minimum requirements. The bottom line is to make sure that you have a good margin between your system's performance and the minimum requirements listed. Always run Windows on more powerful hardware rather than less!

Other hardware—sound cards, network cards, modems, video cards, and so on—may or may not work with Windows. If the device is fairly recent, you can be relatively certain that it was built to work with the newest version of Windows. If it is older, however, you may need to find out who made the hardware and check their website to see if there are drivers available for the version of Windows that you are installing.

There's one more thing to consider when evaluating installation methods. Some methods work only if you're performing a clean installation, not an upgrade.

Understanding Applications

An operating system without any applications installed on it isn't very useful. Even back when Windows 1.0 was released 30 years ago, it came with applications, such as Minesweeper, a file manager, and a simple word processor. In this section, we will explore installing applications for Windows 10/11. However, these concepts can be applied to any operating system.



Exam Objective 1.10, "Given a scenario, install applications according to requirements," requires you to know the various considerations for installing applications. You should understand application requirements, distribution methods, and impact considerations for new applications.

System Requirements for Applications

Before installing an application, you must collect information on the requirements of the application to make sure that your system can satisfy the requirements. If your system cannot satisfy the requirements for the application, you may have to upgrade the system. You can find application requirements on the vendor's website. For example, if you need to install the Microsoft Power BI Desktop application, you would search Microsoft's site to find the system requirements for this particular application. The documentation would reveal several different requirements, such as the following:

Memory The most common requirement that needs to be met before installing an application is the memory requirement. Memory is often displayed in the form of total random access memory (RAM) for the operating system.

CPU The requirement for the *central processing unit (CPU)* required for an application is often displayed in the form of total CPU processing power of the system. It is often expressed in the form of GHz, and other requirements can include the number of physical CPUs, the number of cores, and even the model of processor.

Architecture The architecture of the CPU and the operating system installed needs to match the application requirements for 32-bit or 64-bit applications. The architecture must meet or exceed the requirements of the application. For example, if the application needs a 64-bit architecture, then the CPU and operating system need to support 64-bit operations. However, if the application requires a 32-bit architecture, then a 32-bit or 64-bit CPU and operating system can accommodate the requirement.

Storage The requirement for the storage is often displayed in GB or TB, depending on the application. The storage requirement is often the free space required to install the application. However, some software packages that require the storage of user data also display a separate storage requirement for the data.

Prerequisites A common prerequisite to install an application might be a specific version of the .NET Framework. However, some applications have specific prerequisite requirements, such as a Microsoft SQL Server for data storage.

Display When data is to be displayed in a specific format, the application designer has formatted the application to a minimum display area. This requirement is often in the form of pixels, such as 1024 × 768. Although pixels are the most common display requirement, another type of requirement is a graphics processing unit (GPU). A specific GPU may be required if the application will process video or graphics.

Video Random Access Memory (VRAM) VRAM requirements are typically related to applications that are graphic intensive, such as gaming. VRAM is typically much faster than conventional RAM, because it must be accessed faster to sustain the higher bandwidth requirements for the graphics applications being used.

Built-in Video Cards vs. Dedicated Some video cards that are built into the motherboard of the computer share VRAM with the RAM of the system. These video cards are much slower because they depend on the conventional RAM. These video cards also consume valuable RAM that is typically dedicated to applications. The RAM on the motherboard is divided between the CPU and video card and can be configured through the Basic Input/Output System/Unified Extensible Firmware Interface (BIOS/UEFI) interface. Typically, you can disable these built-in video cards when you install a dedicated video card. Dedicated video cards are always preferable if you are going to use video. The built-in video card is typically an added bonus many vendors include, and typically they are never used.

Licensing When installing an application, the last step is usually licensing the application. The license can come in several different forms, such as a node-locked license tied to a MAC address or drive serial number. A network-based license server allows for concurrent use by several people at once. External hardware tokens license an application with the use of a USB dongle or hardware device that allows the program to run.



Licensing is covered in detail in Volume 2 Chapter 9, “Safety and Environmental Concerns.”

Other In addition to the common requirements, other requirements must be met before an application can be installed. The operating system requirement is one such requirement. Every application will have a minimum operating system on which it can be installed and supported. Other requirements can be for specific hardware, licensing, or drivers.

Additional Application Requirement Considerations

When an application’s requirements are published, they are the bare minimum requirements, not the optimal recommended specifications. The requirements are a conservative estimate so that a potential customer is not scared away with expensive hardware upgrades. When extremely low requirements are advertised, we always recommend calling a presales support person. You should describe how the application will be used and seek their recommendations for optimizing the application’s performance. You will often find that they are double or triple the advertised bare minimum requirements.

Storage is probably the biggest consideration when evaluating the installation of an application that will store user data. It is extremely difficult to gauge how much storage should be set aside for the application for user data. If too much storage is allocated, the precious commodity might go unused for the life of the application. If too little is allocated, you may have to upgrade the storage—or, worse, you could fill the entire drive up and be forced to do it in a panic.

The most common architectures for CPUs are 32-bit and 64-bit. They are noted as x86 for 32-bit and x64 or AMD64 for 64-bit. The 64-bit CPU extension was adopted from the AMD version of their 64-bit CPU. When installing an operating system, the operating system

requirements must be met or exceeded for the CPU architecture type. If you are installing a 64-bit operating system, then you must have at least a 64-bit CPU. If you are installing a 32-bit operating system, you can install it on a 32-bit CPU or 64-bit CPU. However, if you have a 64-bit CPU it always makes sense to install a 64-bit operating system, since it gives you maximum flexibility and you typically can't upgrade to a different architecture.

Operating system compatibility is a consideration for older applications that are still required by organizations. Application compatibility has been incorporated into the Windows operating system since Windows XP. It allows for an application to behave as though the operating system was an older operating system. For example, over the years, Windows has become more restrictive of local permissions on the filesystem. If an application that was written for an older operating system expects to write to a Registry key, application compatibility will allow it to think it's directly writing to the location.

Application Installation

Once the requirements are met for an application and the considerations are investigated, you are ready to install the application. You must now consider how the application will be installed; the number of machines on which the application will be installed will factor into this consideration. This section discusses several different ways that an application can be installed.

Physical Media

If the installation is a one-off installation, then a CD/DVD drive might be your best option. In recent years, applications have even been shipped on Blu-ray media. If you must install the application on several different PCs, then this method may not be the preferred installation method. When installing from optical media, even the fastest optical drive is slow compared to other methods, such as USB flash drives.

ISO Mountable

Although optical discs have been around since the mid-1990s and have been the most popular method of installing applications and operating systems, the optical disc is quickly becoming a relic. When you use a virtualization product such as Hyper-V or VMware Workstation, the optical disc is just too clumsy and slow to use. Mountable ISO images of the physical media has become the new norm. When you need to install an operating system or application, you simply download the media from the vendor, mount the ISO, and install it as if you had a virtual optical drive. ISO sizes will vary from 500 MB to 9.6 GB with normal CD and DVD formats, but Blu-ray discs can be up to 45 GB.

USB Drive

Applications are outgrowing optical media such as CD-ROM and DVD-ROM, so USB drives have become popular. USB drives are faster and bigger than optical media. If a handful of computers require the application, then this might be a better option. However, the disadvantage is that simultaneous installations are limited to the number of flash drives you have

with the application loaded. Another common problem with USB drive installations is that the USB drives are lost from time to time or inadvertently overwritten. For this reason, many application vendors lock their drives so they can't be overwritten and repurposed.

Downloadable Media

Optical and USB media are slowly being replaced, by online market places like Steam for games, and Amazon for software, just to name a few. These market places allow for a downloadable package that can be installed. Some of these market places will even install the application for you from a digital library. These ecosystems are very common with video games and every major video game vendor has their own platform for software delivery.

The software box still has a place in retail, because it satisfies a consumer's need for a tangible product when walking out of the store. However, many software titles sold in a box, will have a digital download code inside.

Image Deployment

Typically, when you have a group of computers that require the same software packages to be installed, it's faster to deploy a software image. A software image is nothing more than a digital image of an operating system with all the software packages installed. Once everything is installed, the image needs to be prepared for deployment with a tool called sysprep. You can then use a tool like Deployment Image Servicing and Management (DISM) to create a digital image of the hard drive. The deployment of the image can be done a number of ways, such as USB, Windows Deployment Services (WDS), Microsoft Deployment Toolkit (MDT), and even DISM. A complete chapter could be dedicated to deployment techniques.

Network Installation

When you need to install an application on many different PCs, a network installation should be your first choice. The application is typically uploaded to a file share by the administrator, and then the file share is set to read-only access for the user who is performing the installation. Depending on the speed of the network, this could be the fastest method to install an application on several different PCs simultaneously.

There are several different methods for deploying an application over the network. Each method depends on the number and location of the computers.

User-Initiated Installation

These installations are used when the administrator of the PC will start the installation manually. This method is preferable when the administrator is expected to answer specific questions during the installation, such as where to install the application.

The installation of any application generally requires that the user be an administrator of the operating system or have elevated privileges to install an application. If user-initiated installation is chosen as the method for deploying an application, you must be sure that the user account has the appropriate permissions to install the application.

Automated Installation

Two different types of automated installations can be employed by the administrator: push installations and pull installations. Either installation type is used when the conformity of the installation is required.

Automated installation products, such as Microsoft Endpoint Configuration Manager (MECM), can deploy applications to multiple PCs with a push installation. Microsoft MECM, is considered the Swiss Army knife of installation and reporting services. Each client in the network will have the MECM agent installed prior to the push installation of the application. The agent will then be responsible for reporting on the current operating system. This allows you to collect information for determining if you can satisfy the requirements of the application. The agent can then be used as the contact point for the push installation of the application.

Group Policy can also be used to automate the installation of applications. The Group Policy method is a pull-based installation method, where the client will pull the application from the network share. This method contains no agent, so reporting is not available for client resources, installation requirements, or installation status. The benefit is that it requires very little infrastructure if you have Microsoft Active Directory (AD) installed already.

Regardless of the installation type, push- or pull-based, most automated installation methods do not require the user logged in to be an administrator. MECM can be configured to use a system account that has elevated privileges in the target operating system to install the application. Group Policy pull-based installations can also be configured to use a system account that has elevated privileges to install the application.



Push installations are installations in which the administrator can push an application out to clients. Pull installations are installations in which the client pulls, or queries, for the application to be installed. Group Policy application deployment is a pull type of installation, since the client is responsible for pulling the application and installing it. MECM is a push type installation method. For more information about MECM, visit:

<https://learn.microsoft.com/en-us/mem/configmgr/core/understand/configuration-manager-faq>.

Security Considerations

When an application is installed in the operating system, the overall security of the operating system can potentially be compromised. Of course, it is not the intent of the application to weaken the security of the operating system, but vulnerabilities in the application can exist. Application developers realize that their code is not perfect and that vulnerabilities can exist, so many applications now include self-updating capabilities.

Applications weakening the security of the local device should not be your only concern. Applications can actually weaken the security of the entire network, especially if they are not updated frequently. Applications that operate over the network with client-server

functionality pose the largest security risks, since they allow for remote exploitation of vulnerabilities. One measure that can be employed with network applications that are not updated frequently is to firewall the services. There are many other active measures, which we will discuss in further detail in Volume 2 Chapter 6, “Securing Operating Systems.”

Other Considerations for New Applications

Once the requirements of the application are satisfied, the security of the application is evaluated, and an installation method is chosen; several other considerations should be reviewed before installing a new application. These considerations must be evaluated to ensure that the installation of the application does not impose unintended problems on the organization and network.

Impact to Device Although the impact to the device should be minimal if the application requirements are satisfied, the impact should still be considered. For example, if you install an application that constantly logs, you could find that the hard drive will fill up quickly and other applications could run more slowly due to the drive activity. Although this is an extreme example, other impacts may be more subtle and manifest over a longer period of time.

Impact to Network The impact to the network should also be considered. The bandwidth the application needs to operate, and consequently uses, should be evaluated. Although an application at first glance might not appear to use a lot of bandwidth, as more demand is placed on the application, bandwidth will increase. For example, if you installed a network video recorder (NVR), a few cameras might not use too much network. However, as more and more cameras are added, bandwidth usage can sharply increase. If the bandwidth traverses a wide area network (WAN), the application could quickly use all the available bandwidth.

Impact to Operations Depending on the type of application being installed, the operations of the operating system, or the information system the operating system serves, can be negatively impacted. Therefore, the impact to operations should always be considered before installing an application. For example, if you were to install an Office package on a point-of-sale (POS) system, you could negatively affect the ability of the system to process sales.

Impact to Business Some applications can have an impact on the business or organization itself. Therefore, the impact on day-to-day business should be evaluated. Examples of applications that could affect business are human resource (HR), customer relationship management (CRM), and sales systems, just to name a few. If these applications were to be changed, they could impact the organization.

Cloud-Based Productivity Tools

The cloud was a buzzword 10 years ago and slowly gained traction, as administrators held tightly onto their on-premises servers with skepticism. Since the COVID 19 pandemic of 2020, the cloud is no longer a buzzword. During the pandemic every administrator who had not embraced the cloud found themselves rapidly deploying cloud-based applications.

Today, every system put in place should be considered for the cloud—not because we expect another pandemic, but because the work force has forever changed! Pandora’s box is open, and organizations have found out that people are actually productive when they work from home (WFH). Working from home also benefits the employee and affords a greater degree of life-balance. No more worrying that the kids will get off the bus and come home to an empty house, since you are working from your home office. Many jobs today are considered hybrid, meaning the employee works from home a few days a week and also comes into the office a few days a week. Some jobs are completely remote!

All of this means that at some point your organization needs to embrace cloud-based productivity tools. The following are the most common productivity tools you may need to support. We will focus primarily on Microsoft 365 cloud-based tools in this section, mainly because they are the most common. However, keep in mind that they are not the only tools out there in the marketplace.



Exam Objective 1.11, “Given a scenario, install and configure cloud-based productivity tools,” requires you to know the various cloud-based productivity tools in the following section.

Cloud-Based vs. On-Premises

The great debate of cloud-based versus on-premises productivity software centers on flexibility, cost, and control. The flexibility of on-premises systems is only as good as your Internet connection and how you have deployed the system. If you’ve deployed the email server to a single server, it is possible that demand from the outside of the organization could impact users inside the organization. This could be normal usage or a threat actor impacting the email system. Alternately, if you have a poor Internet connection, users outside could have a poor experience while internal users are fine.

Cost is another consideration in the debate between cloud-based and on-premises systems. Typically, when you set up an on-premises server there is a one-time cost for the infrastructure, such as the server and the software. Keeping in mind that, although your organization might use virtualization, there is still a shared cost for the CPU, RAM, and disk space. After the initial costs, there might be a support contract or licensing that is yearly.

When we purchase a cloud-based service, we are renting the infrastructure. Typically cloud-based services are charged per-user/per-month, with a discount for committing to a full year and paying up front. The positive is you don’t need to purchase servers or tap into precious virtualization resources. Also, most cloud-based systems include basic support costs. The negative is that you need to pay a larger amount each year you use the service.

The last consideration is control. When we use a cloud-based service we give up control to the vendor; this could be good or bad. If you need complete control over the version of the system and how the system is configured, then cloud-based might be a bad decision. However, updates and new versions means value for the customer. Another big benefit is security; when applications become cloud-based, you no longer need to worry about patching the software. The vendor will typically have a continuous patch cycle and offer release information on new

features in scheduled releases. Organization can realize value from cloud-based systems, when they simply don't have the expertise on hand to maintain the systems.

Cloud-Based vs. Traditional Software

Cloud-based software is typically delivered via a standard web browser and requires very little configuration when deployed in this fashion. The storage and processing are all performed on the remote servers—only the user interface is rendered on the client. The only processing on the client is the sorting or presentation of data.

Traditional software typically needs to be installed on the operating system, and the install is specific to the application. Storage and processing are typically all done on the local operating system. In some cases, you can point the application to a centralized server for the storage of information.

Cloud-based software that is only used via a web browser can be cumbersome. Therefore, today's cloud apps allow a traditional install of the application that points to the cloud for storage. The cloud application might also allow a web interface. This is the case with the Microsoft 365 suite of products; you can install the Microsoft 365 applications or you can use them through the `Office.com` web portal. Keep in mind that the applications installed on the operating systems typically have more advanced features than their web-based browser counterpart.

This marriage of web-based cloud app and traditional installation gives you tremendous flexibility. You can utilize a Google Chromebook to do day-to-day tasks via the web browser, then a Windows computer for specific tasks that require the full version. Since the storage and format of the document are the same, you can literally pick right back up where you left off, regardless of which platform you are on. This also means you can use the Microsoft 365 applications on computers where you might not have administrative rights to install software and just need access to a web browser.

Cloud-Enabled Applications

Connectivity is a consideration when using cloud-based applications. You need access to an Internet connection to use them. This is where a traditional installation or a cloud-enabled application is a benefit. In the following sections, we will explore two cloud-enabled applications that are installed locally on your computer. They both allow offline use and then communicate with the cloud when Internet is restored.

Exchange Online and Outlook


Microsoft Exchange has been the foundation of the Microsoft email and collaboration services since its initial release in 1993. The on-premises Exchange Server offered organizations a robust solution for email, calendaring, contact management, and overall communications. However, with the rise of cloud computing, this led to the launch of Exchange Online in 2008 as part of Microsoft's Business Productivity Online Suite (BPOS) and later integrated into Office 365 (now Microsoft 365). Unlike traditional on-premises Exchange Server, Exchange Online operates entirely in the cloud. This eliminates the need for complex local

infrastructure while providing seamless updates, scalability, and advanced features like AI-driven insights, enhanced security, and compliance tools.

Microsoft Outlook was introduced in 1997 as part of the Office 97 suite. It quickly became the keystone of personal and professional email. Initially designed as a desktop application, Outlook provided users with a unified platform for email, calendars, tasks, and contact management, primarily connecting to on-premises servers like Microsoft Exchange. Over time, Outlook evolved with features such as integration with mobile devices, advanced search, and shared calendaring, just to name a few. Today, paired with Exchange Online, Outlook has transcended its desktop roots to offer a robust cloud-enabled application.

Just like its initial debut with the install of Office 97, today it is installed as an integral part of the Microsoft 365 Apps suite. The Microsoft 365 Apps suite is even preinstalled in Windows 11 by default; all you need to do is sign in with your Microsoft Account credentials. If you need the installation of Microsoft 365 Apps, you can visit www.microsoft365.com and click Install And More, then select Install Microsoft 365 Apps. This will take you to the familiar dialog shown in Figure 1.2. The executable that downloads will install the latest Microsoft 365 Apps (Office) applications of Word, Excel, PowerPoint, OneNote, Access, Publisher, and of course Outlook.

FIGURE 1.2 Microsoft 365 Apps (Office) download




When you first launch Outlook, it will ask you to sign in. After entering your email address, it will check DNS for an autodiscover record that resolves to the Exchange Online service. Outlook will then set up your profile pointed to the Exchange Online service.

Outlook will be configured in Cached Exchange Mode by default if you use an Exchange Server or Microsoft 365 account. Cached Exchange Mode allows you to work offline by downloading the last 12 months of email to the computer. When you search, it will search this first 12 months, then ask you if you want to expand the search to online. When you create new emails and don't have an Internet connection, the emails will sit in the outgoing mailbox until you connect to the Internet and launch Outlook again.

You can turn Cached Exchange Mode on or off manually, or even vary the length of time email is cached with the following procedure:

1. In the Outlook client, click File > Account Settings > Account Settings.
2. Click Exchange or Microsoft 365, and then click Change.
3. Under Offline Settings, select Use Cached Exchange Mode to turn it on or off, as shown in Figure 1.3.

FIGURE 1.3 Cached Exchange Mode



4. Move the slider to lower or increase the length of time emails are cached.
5. Click Next, then Done to save the changes.



The autodiscover record is created by the administrator when the Exchange Online server is initially configured for an organization. If you are using a personal Microsoft 365 account, <http://autodiscover.outlook.com> points to Microsoft's internal Outlook.com infrastructure.

Microsoft OneDrive

Microsoft OneDrive was launched in 2007 as Windows Live SkyDrive. It began as a simple online storage service for users to save and share files. Over time, it evolved into a cloud storage platform and was rebranded as Microsoft OneDrive in 2014.

OneDrive is integrated tightly with Microsoft 365 and in Windows 11, OneDrive is an integral part of the File Explorer interface. A dedicated OneDrive section appears in the navigation pane, allowing users to access their cloud-stored files alongside local ones. Additionally, integration with Windows 11's notification system ensures users are alerted about file activity, such as completed syncs, sharing updates, or storage limit warnings.

OneDrive provides users with effortless access to their files from any device, along with advanced collaboration tools and automatic backup capabilities. A standout feature of OneDrive is its synchronization functionality, which enables users to keep their files updated across devices and the cloud.


The Files On-Demand feature is also enabled by default, displaying cloud files as placeholders until opened, which conserves local disk space. OneDrive also includes built-in sharing options, encryption, and integration with tools like Word, Excel, and Teams, making it a vital part of both personal and professional file management.

During installation or setup, the user signs in with their Microsoft or organization account to link their OneDrive storage. By default, key system folders like Desktop, Documents, and Pictures are offered for automatic backup and sync to OneDrive, ensuring critical files are protected without additional configuration. Through the OneDrive app, users can select specific folders to sync, ensuring that their essential files are always current across devices.

To change the current folders syncing to OneDrive, use the following procedure:

1. Click the OneDrive cloud icon in the Taskbar.
2. Select the gear icon in the upper-right corner.
3. Choose Settings from the drop-down menu.
4. Choose Account from the left menu.

The Choose Folders dialog will pop up, where you can choose which files are synced with OneDrive, as shown in Figure 1.4.

FIGURE 1.4 OneDrive Choose Folders

Microsoft OneDrive is not the only cloud-enabled storage sync tool available. Google Drive is another popular option, and if you are on a Mac, then iCloud is the preferred method. There are also third-party solutions like Carbonite Backup, DropBox, and Backblaze—these are just a few of the many different vendor offerings.

Collaboration Tools

Cloud-enabled collaboration tools have revolutionized the way teams work together, offering seamless integration and real-time collaboration across various platforms. The Microsoft 365 suite, which includes tools like Microsoft Teams, SharePoint, and OneDrive, provides a comprehensive solution for businesses to communicate, share files, and collaborate on projects from anywhere in the world. These tools are designed to enhance productivity and streamline workflows, making it easier for teams to stay connected and work efficiently.

Google Workspace, formerly known as G Suite, offers a robust set of applications such as Google Drive, Google Docs, and Google Meet. These tools enable users to create, edit, and share documents in real time, facilitating smooth collaboration and communication.

Both Microsoft 365 and Google Workspace are equipped with advanced security features, ensuring that data is protected while enabling teams to collaborate effectively in a cloud-based environment. Microsoft 365 Productivity Apps are typically bundled with Word, Excel, PowerPoint, Outlook, Access, Teams, and of course OneDrive to back up and store all your documents. Google Workspace Apps are typically bundled with Docs, Sheets, Slides, Gmail, Meet, Calendar, Chat, and of course Google Drive to back up and store everything. Here are some of the apps you might encounter:

Word Processing Tools Word processing applications such as Microsoft Word, and Google Docs offer robust platforms for creating, editing, and sharing documents. These tools support real-time collaboration, enabling multiple users to work on the same document simultaneously. With features like Track Changes, comments, and various formatting options, users can efficiently produce and refine their written content. The cloud integration ensures that documents are automatically saved and accessible from any device.

Spreadsheets Cloud-enabled spreadsheet applications like Microsoft Excel and Google Sheets have transformed data management and analysis. These tools allow multiple users to collaborate in real time, making it easy to update and share data instantly. With features such as pivot tables, charts, and various formulas, users can efficiently organize and interpret large data sets. These functionalities enable users to perform detailed data analysis, create visual representations of data, and generate insightful reports. The ability to use complex formulas and functions further enhances the analytical capabilities of these tools, making them indispensable for data-driven decision-making. The cloud integration ensures that all changes are saved automatically and accessible from any device, enhancing productivity and collaboration.

Videoconferencing Videoconferencing tools such as Microsoft Teams and Google Meet have become essential for remote communication and collaboration. These platforms offer high-quality video and audio calls, screen sharing, and virtual meeting rooms, enabling teams to connect and collaborate from anywhere in the world. With features like chat, file sharing, and integration with other productivity tools, videoconferencing applications facilitate seamless communication and efficient teamwork. Microsoft Teams and Google Meet are the two dominant videoconferencing solutions in enterprises today. However, there are many other third-party products, such as Zoom, GoTo Meeting, and Cisco Webex, just to name a few popular apps; there are many others.

Presentation Tools Cloud-enabled presentation tools like Microsoft PowerPoint and Google Slides provide powerful platforms for creating and sharing engaging presentations. These tools support real-time collaboration, allowing multiple users to work on the same presentation simultaneously. With a wide range of templates, design options, and multimedia integration, users can create visually appealing and informative presentations. The multimedia integration capabilities of these tools allow users to embed videos, images, and audio clips, enhancing the overall impact of their presentations. Additionally, features like animations and transitions can be used to make the content

more engaging and interactive. The cloud storage ensures that presentations are always accessible and up-to-date.

Instant Messaging Instant messaging tools like Microsoft Teams and Google Chat provide quick and efficient communication channels for teams. These platforms support real-time text messaging, file sharing, and integration with other productivity tools, making it easy for team members to stay connected and collaborate. With features like group chats, direct messages, and notifications, instant messaging applications enhance communication and streamline workflows. There are many other instant messaging cloud-enabled apps, but these are the two most popular in enterprises today.

Hybrid Identity

Active Directory (AD) is a directory service developed by Microsoft for Microsoft Windows domain networks. It provides a centralized system for authentication, authorization, and user management, and these are just a few capabilities. AD stores information about objects on the network and makes this information easy for administrators and users to find and use. AD is typically used for on-premises networks in which you need to authenticate to a domain controller.

Entra ID, formerly known as Azure Active Directory (Azure AD), is Microsoft's cloud-based identity and access management (IAM) service. Entra ID also allows for authentication, authorization, and user management, in addition to many other functions. Some of the security features Entra ID supports are single sign-on (SSO), multifactor authentication (MFA), and conditional access policies to enhance security and streamline user access, and that's just scratching the surface. Entra ID is considered an identity provider (IdP), and it integrates with any cloud application that supports SSO. Entra ID is also required if you are using any Microsoft 365 platform apps, such as Exchange Online or OneDrive, just to name a few.

Both AD and Entra ID serve a similar function and support identities differently, depending if you are authenticating on a domain-joined computer or signing into a cloud-enabled application. Typically, when an organization decides to expand to the Microsoft 365 cloud platform, it is necessary to support a hybrid identity. You can achieve a hybrid identity two main ways; with Entra ID Connect Sync or Entra ID Cloud Sync. However, the basic idea is synchronizing AD identities to Entra ID, and both tools achieve this. Therefore, when a domain user is created, a cloud-based identity will also be created in Entra ID to support cloud-based services.

Entra ID Connect Sync performs a one-way synchronization of users, groups, and computers in AD to Entra ID. It can be installed on only one server, and it's typically installed on the domain controller. It requires dedicated credentials and synchronizes every 30 minutes. Entra ID Connect Sync has been around since the introduction of Azure.

A newer approach is installing and configuring Entra Cloud Sync. Entra ID Cloud Sync is an agent-based installation that is installed on your domain controller. It does not require a service account and can be installed on more than one domain controller. It syncs every 2 minutes, and it can even sync cloud-based groups back to the AD domain. Entra ID Cloud Sync is configured and controlled through the cloud, as opposed to Entra ID Connect Sync,

which is administered on the server it is installed on. Obviously Entra ID Cloud Sync has a ton of features compared to the previous version and should be your preference.

Both Entra ID Connect Sync and Entra ID Cloud Sync both support password writeback. This feature enables you to change your password on the Office.com platform and the password will write back to the AD domain.

Cloud Licensing

When a user, also known as an identity, is created in the cloud, it is created because we want to allow the user access to cloud-based services. Cloud licensing for the Microsoft 365 platform provides a flexible and scalable way for organizations to manage user access to various Microsoft services. With cloud licensing, administrators can assign licenses to individual user accounts or groups of users, enabling them to use applications such as Microsoft Teams, Exchange Online, and OneDrive, just to name a few.

Microsoft offers licensing to organizations in the form of tiers for the Microsoft 365 platform. Each tier will add additional functionality to Exchange, security, compliance, and Entra ID, among others. Each tier is also priced incrementally based on the features, and they are charged per user/per month. The two major commercial tiers for Microsoft 365 are E3 and E5. Academic institutions have three tiers—A1, A3, and A5—with the lowest tier typically used for students.

Along with the Microsoft 365 platform, you can purchase additional licensing for other Office products. There are many other add-ons, such as telephony and larger storage for Exchange, among other add-ons.

To associate a license, an administrator would log into the Microsoft 365 admin center, navigate to the Users section, select the user account, and then assign the appropriate license from the available options. Additionally, an administrator can assign a group to a license; this is called group-based licensing. The benefit is that when a user is added to the specific group, they automatically are assigned a license.

Introduction to Windows

Microsoft released Windows 10 on July 29, 2015 and offered a free upgrade from Windows 7, Windows 8, and Windows 8.1. The free upgrade was to be limited to the first year of the Windows 10 release and expired on July 29, 2016. However, it's been reported that Microsoft still allows the upgrade and activation of Windows 10 for free, but your luck may vary. For these reasons, Windows 10 was rapidly adopted. Microsoft released Windows 11 on October 5, 2021, and it too is offered as a free upgrade for any Windows 10 device that meets the minimum requirements. Windows 10 is scheduled for end-of-life October 14, 2025, and Windows 11 will become the current version.



Exam Objective 1.3, “Compare and contrast basic features of Microsoft Windows editions,” requires you to know the difference between the Windows 10 and 11 versions. Be sure to understand the feature differences.

Windows Editions

Windows 11 has had a total of 10 editions released. Only two of the 10 editions were made available in the retail channel: Windows 11 Home and Windows 11 Pro. Microsoft has made it easy for the retail consumer to pick the right edition of Windows 11, narrowing down retail choices to the two editions. Technically, there is a third retail edition called Windows 11 Pro for Workstations, which is preinstalled on high-end hardware for *high-performance computing (HPC)* requirements and allows up to 4 CPUs and 6 TB of RAM.

The editions available for Microsoft volume licensing options are Windows 11 Enterprise, Windows 11 Enterprise LTSC (Long-Term Servicing Channel), Windows 11 Education, and Windows 11 Pro Education. Windows 11 volume license editions include features such as AppLocker, BranchCache, and DirectAccess, just to name a few. Windows 11 Enterprise LTSC is an edition that is released every 2–3 years and is supported for 10 years after its initial release. Windows 11 LTSC receives normal Windows updates for security but does not receive feature upgrades. The *Microsoft Store* and bundled apps are also omitted from the Windows 11 Enterprise LTSC edition. Both Windows 11 Enterprise and Windows 11 Education editions have the same features over and above Windows 11 Pro and Windows 11 Home editions. Windows 11 Education editions are made available only to academic institutions and K–12 schools.

Although there are 10 different editions of Windows, the 220-1202 exam focuses on the following four Windows 10 editions: Home, Pro, Pro for Workstations, and Enterprise. In addition, the 220-1202 exam focuses on the three editions of Windows 11: Home, Pro, and Enterprise. Although Windows 11 Pro for Workstations is not an exam objective, we included it for completeness.

Windows 10 Pro and Windows 11 Home This edition is sold as a retail product, and it's preinstalled by the hardware vendor. When the operating system is preinstalled, it is referred to as an original equipment manufacturer (OEM) operating system. Windows 10/11 Home also lacks the key feature of joining a domain that can only be found in Windows 10/11 Pro or Enterprise.

Windows 10 Pro and Windows 11 Pro (Professional) This is the premier retail edition of the Windows 10/11 operating system. It is the most common OEM operating system because it allows users to join corporate domains and becomes value-added for the hardware vendor.

Windows 10 Pro for Workstations and Windows 11 Pro for Workstations This edition is used primarily for high-end computer workstations that require more than two CPUs and that need to address more than 2 TB of memory. Windows 10/11 Pro for Workstations supports up to four CPUs and 6 TB of memory. In addition, it supports Remote Direct Memory Access and Non-Volatile Dual In-line Memory Modules (NVDIMM).

Windows 10 Enterprise and Windows 11 Enterprise This edition supports all features of Windows 10/11 Pro and can be found in mid- to large-sized organizations. It also

supports features generally found in mid- to large-sized organizations. Windows 10/11 Enterprise can be purchased with a volume license subscription from Microsoft or an E5 Microsoft 365 license.



More information on Windows 11 editions can be found here:

[www.microsoft.com/en-us/windows/
compare-windows-11-home-vs-pro-versions#tabs1-2](https://www.microsoft.com/en-us/windows/compare-windows-11-home-vs-pro-versions#tabs1-2)

There are 32-bit and 64-bit versions available for each of the Windows 10 editions listed except Windows 10 Pro for Workstations, since it is specifically used for high-performance computing. With the introduction of Windows 11, Microsoft only offers Windows 11 in a 64-bit version.

N Version

The Windows N version is a special edition of the Windows operating system designed to comply with European Union (EU) regulations. Unlike standard Windows editions, Windows N versions do not include built-in multimedia features such as Windows Media Player, Groove Music, Movies & TV, Voice Recorder, or Skype. This decision was made to provide users with the freedom to choose and install their preferred third-party media applications, leveling the playing field for other software developers. The absence of these features means that users may need to install the Media Feature Pack to restore multimedia functionality. This edition is particularly useful for those who prefer a more customizable operating system or have specific requirements for their multimedia software.



For Exam Objective 1.2, “Given a scenario, perform OS installations and upgrades in a diverse environment,” be sure to understand upgrade considerations for various edition upgrades and operating system upgrades covered in the following section.

Upgrade Paths

If you want to perform an upgrade of Windows 11 instead of a clean installation, there are a few rules. When you upgrade from Windows 10 Home to Windows 11 Pro, you’ll need to provide the new activation key for the higher edition. You cannot upgrade to Windows 10 to a lower version. Although you can upgrade the edition when performing an operating system upgrade, it is recommended that you use like-to-like editions of Windows. A like-to-like edition is an edition with similar functionality, such as Windows 10 Home to Windows 11 Home, or Windows 10 Pro to Windows 11 Pro.

Upgrading from Windows 10 to Windows 11 requires meeting several key prerequisites to ensure a smooth transition. First and foremost, your device must have a compatible 64-bit processor with at least 1 GHz clock speed and two or more cores. Additionally, a minimum of 4 GB of RAM and 64 GB of storage is essential. The system firmware must support Unified Extensible Firmware Interface (UEFI) and the ability for Secure Boot. The motherboard must have a Trusted Platform Module (TPM) version 2.0 enabled for enhanced security. A DirectX 12-compatible graphics card with a Windows Display Driver Model 2.0 (WDDM) driver is also necessary. Ensuring these requirements are met will help facilitate a seamless upgrade to Windows 11, providing access to its advanced features and improved performance.

If you want to retain settings and applications, an in-place upgrade is required in which you upgrade the current operating system to the desired version.

Another major restriction is that you cannot switch architecture during an in-place upgrade. If you have a 32-bit installation of Windows 10, then you will need to perform a clean install of Windows 11. Windows 11 supports only the 64-bit architecture, not the 32-bit architecture. You can switch architectures only when performing a clean installation, in which the old operating system is overwritten. This means that you must back up and restore settings, files, and applications.

It is always recommended to have a backup of your installation before you begin upgrading. It is equally important to check the application compatibility before upgrading. You should always check with the vendor of the application prior to upgrading to make sure they support Windows 11. An upgrade of the application may be required, or a completely new version of the application may be needed.

In addition to checking application compatibility prior to upgrading, you should check the hardware compatibility. Most of the time, a simple upgraded driver is required; however, sometimes you will find that the hardware is not compatible and that new hardware is required. In rare instances, the PC hardware is just not compatible and a new PC is required. This is usually the case with laptops, since components are not easily upgradable or upgrades are impossible. So do your homework prior to upgrading and know what to expect after the upgrade process is complete.

Windows Features

Every Windows operating system edition has a set of features that are bundled into the edition being purchased. The difference between the four main editions of the operating system—Home, Pro, Pro for Workstations, and Enterprise—is how the edition is purchased and its accompanying features. Windows 10/11 Home and Windows 10/11 Pro are both available through retail channels and can be purchased off the shelf. Windows 10/11 Pro for Workstations is typically preinstalled on high-performance workstations. Windows 10/11 Enterprise edition is unavailable without a volume license agreement from Microsoft or E5 Microsoft 365 subscription. Table 1.2 compares the features of the four main editions of Windows 10/11.

TABLE 1.2 Windows 10/11 features and editions

Edition	Maximum RAM supported	Maximum physical CPUs supported (multiple cores)	Notes
Home	128 GB	2	Lacks support for Remote Desktop (client only), BitLocker, Windows To Go, Hyper-V, joining to a domain, and participating in Group Policy. This edition is strictly for consumer use.
Pro	2TB	2	Can join a Windows server domain; includes Remote Desktop Server, BitLocker, Windows To Go, Hyper-V, and participating in Group Policy.
Pro for Workstations	6TB	4	Includes all features of the Pro edition, with support for 6TB of RAM and up to 4 physical CPUs.
Enterprise	6TB	2	Includes BitLocker, support for domain joining and Group Policy, DirectAccess, AppLocker, and BranchCache. This edition is available only through a volume license subscription.


The following is a list of features introduced and associated with the Windows 10/11 operating system that you should know for the exam, along with a brief description of each:

Domain vs. Workgroup Windows 10/11 Pro edition offers a significant advantage for businesses through its support for domain joining. Domain joining is a feature that allows computers to be part of a centralized network managed by Active Directory. This setup provides enhanced security, centralized management, and streamlined access to network resources, making it ideal for mid to larger organizations. In contrast, a workgroup is a simpler network model where each computer is managed independently, suitable for smaller networks or home environments. While workgroups offer ease of setup and maintenance, they lack the robust security and administrative controls available in a domain environment. Domain joining is not a feature in Windows 10/11 Home edition.


Remote Desktop Protocol Windows 10/11 Pro edition includes the Remote Desktop Protocol (RDP), a powerful feature that allows administrators to remotely access and control the operating system. RDP provides a secure and seamless connection, enabling administrators to run applications, manage files, and perform tasks as if they were physically present at the computer. For example, a network administrator can use RDP to troubleshoot issues on a workstation located in a different office. The Windows 10/11 Home edition includes the Remote Desktop client, but the operating system does not permit remote access connections inbound.

Microsoft Edge The Microsoft Edge web browser was developed as a lightweight web browser, as shown in Figure 1.5. Microsoft Edge is the successor to Internet Explorer and has already replaced Internet Explorer as the default web browser on Windows 10/11. The original version of Edge was released on January 15, 2020, and it was redesigned based upon the Chromium source code. Microsoft Edge does not support ActiveX and browser helper objects. Internet Explorer 11 reached end-of-life (EOL) on June 15, 2022. Internet Explorer 11 is no longer installed in Windows 11. You can still use the IE Mode in Microsoft Edge to provide backward compatibility. The Microsoft Edge browser can be seen in Figure 1.5.

FIGURE 1.5 The Microsoft Edge web browser





BitLocker BitLocker allows you to use drive encryption to protect files, including those needed for startup and login, as shown in Figure 1.6. This feature is available in all editions of Windows 10/11, except for Windows 10/11 Home. For removable drives, BitLocker To Go provides the same encryption technology to help prevent unauthorized access to the files stored on them. BitLocker is covered in Volume 2 Chapter 6.

FIGURE 1.6 The BitLocker Control Panel applet

Group Policy Group Policy is a powerful feature in Windows 11 that allows administrators to manage and configure operating system settings for the local computer. When used with an Active Directory domain, it provides centralized control over user and computer configurations, enhancing security and efficiency over multiple computers. The `gpedit.msc` tool, also known as the Group Policy Editor, is the interface used to create, modify, and apply these policies. This tool is available in all Windows 10/11 editions, except for the Windows 10/11 Home edition.

Desktop Styles Windows offers a variety of desktop styles that allow you to personalize your experience. From the sleek and modern look of Windows 11 to the classic feel of Windows 10, each version provides unique customization options. Users can choose from a range of themes that change the background, window colors, and sounds. You can even download new themes from the Microsoft Store. Additionally, Windows 11 introduces three different modes: Dark, Light, and Custom, each with its own color customization options for the Start menu, Taskbar, and Action Center, as shown in Figure 1.7.

Microsoft Defender Windows 10/11 includes an antivirus, antimalware utility that is built into the operating system. You can always provide your own antivirus or anti-malware, but Microsoft Defender allows for protection from the time that Windows is installed. In addition to antivirus/antimalware, Microsoft Defender can protect against ransomware, as shown in Figure 1.8.


FIGURE 1.7 WindowsTheme settings**FIGURE 1.8** Microsoft Defender



Microsoft Defender is covered in detail in Volume 2 Chapter 6.


Settings App The Settings app is similar to Control Panel in functionality, because you can control many aspects of the operating system. However, any new feature being added to the Windows 10/11 operating system is added to the Settings app in lieu of the creation of new Control Panel applets. At this point, it looks like Control Panel will become a legacy component in future versions of Windows. You can access the Settings app by clicking the Start Menu and selecting the gear on the left-hand side. You can also right-click the Start menu and select Settings from the context menu. The Settings app is shown in Figure 1.9.

FIGURE 1.9 The Settings app in Windows 11




Pinning to the Taskbar The context menu for each application allows you to choose whether you want to pin (add) it to the Start menu, pin it to the Taskbar, or remove it if it is already there. You can also uninstall an application with the context menu, as shown in Figure 1.10.

FIGURE 1.10 An application context menu



Snap Assist Windows Snap was first introduced with Windows 7. Snap allows for side-by-side automatic windowing of windows on the screen. All you have to do is move a window over to a side, and the window will resize to half of the screen. Windows 10 introduced Snap Assist, which allows for up to four quadrants of the screen to automatically be adjusted using snapping. When a window is snapped over to a side, a preview in the blank area will show a preview of windows so that you can choose the next window to snap. Windows 11 added the functionality of dragging the window to the top drop-down bar to automatically place the window in a division, after the first one has been snapped in place, as shown in Figure 1.11.

FIGURE 1.11 Windows Snap Assist

Windows Hello Windows Hello was first introduced as a new feature of Windows 10 that allows the user to sign into Windows with biometrics. Windows Hello uses the built-in camera on the laptop or PC so that the device can be unlocked with facial recognition. On devices with a fingerprint reader, you can also use your fingerprint to sign into Windows.



Windows features are added with every version/update of the Windows operating system. You can view information about upcoming features that Microsoft is working on with the Windows Insider program. For more information, visit <https://insider.windows.com/en-us>.

The Windows Interface

If you've worked with older versions of Windows (such as Windows 7), you'll notice that it looks similar to the current Windows interface. While there are some differences, most of the basic tasks are accomplished in an almost identical fashion on everything from a Windows 95 workstation computer on up. Also, although the tools that are used often vary between the different OSs, the way that you use those tools remains remarkably consistent across platforms.

We will begin with an overview of the common elements of the Windows GUI. We will then look at some tasks that are similar across Windows operating systems. You are encouraged to follow along by exploring each of the elements as they are discussed.



As you follow along, you may notice that there are numerous icons and options that we do not mention. Quite honestly, there are too many to cover, and they're beyond the scope of this chapter. For now, simply ignore them or browse through them on your own, and then return to the text.


The Desktop

The *Desktop* is the virtual desk on which all your other programs and utilities run. By default, it contains the Start menu, the Taskbar, and a number of icons. The Desktop can contain additional elements, such as shortcuts or links to web page content. Because the

Desktop is the foundation on which everything else sits, the way that the Desktop is configured can have a major effect on how the GUI looks and how convenient it is for users. When you click the lower-left corner of the Desktop, the Start menu appears. (Right-clicking the Windows icon in Windows 8.1 and above displays a set of operating system functions.)

You can change the background patterns, screen saver, color scheme, and size of elements on the Desktop by right-clicking any area of the Desktop that doesn't contain an icon. The menu that appears, similar to the one shown for Windows 11 in Figure 1.12, allows you to do several things, such as create new Desktop items, change how your icons are arranged, or personalize your desktop.

FIGURE 1.12 The Windows 11 Desktop context menu




The Three Clicks in Windows

When it comes to interacting with a mouse in Windows, there are three possibilities:

- **Primary Mouse Click**—A single click (typically the left mouse button) is used to select an object or place a cursor.
- **Double-Click**—Two primary mouse clicks in quick succession are used to open a program through an icon or for other application-specific functions.
- **Secondary Mouse Click (or Alternate Click)**—Most mice have two buttons. Clicking once on the secondary button (usually the one on the right, although that can be modified) is interpreted differently from a left mouse click. Generally, this click displays a context-sensitive menu in Windows from which you can perform tasks or view object properties.

When you right-click the Desktop in Windows 11 and choose Personalize, you will see the Personalization Settings screen, as shown in Figure 1.13.

FIGURE 1.13 The Windows 11 Personalization setting screen

With the rapid adoption of Windows 11, this book will cover the most current version of Windows 11 (24H2), since the CompTIA objectives focus on Windows 11. However, every prior operating system Microsoft has produced has similar settings for personalization. We will cover the main ones in the Display Settings window for Windows 11.

Background The Background section is used to select a picture to display on the Desktop. In addition to setting a picture, you can select a group of pictures to display as a slideshow or display a static color.

Colors The Color section allows you to select a color scheme for the Start menu or to change the color or size of other Desktop elements.

Themes The Themes section allows you to select a theme to customize the look and feel of your machine quickly. Selecting a theme sets several items at once, such as a picture to display on the Desktop, the look of icons, the sounds to use, and so on. All these options can also be selected individually through the other personalization tabs.

Dynamic Lighting Dynamic lighting in Windows 11 settings allows users to manage RGB lighting for their devices directly from the operating system, eliminating the need for third-party applications.

Lock Screen Windows 8 and above allows you to easily change the lock screen wallpaper as well as the applications that are displayed on the lock screen.

Text Input In Windows 11, Text Input settings offer advanced options to enhance your typing experience. You can enable features such as autocorrect misspelled words, highlight misspelled words, and show text suggestions as you type. These settings help improve typing accuracy and speed, making it easier to compose messages and documents. Additionally, you can customize the touch keyboard settings to suit your preferences and enhance your overall typing experience.

Screen Saver Although the screen saver is not in the Personalization menu, it's worth covering. The screen saver can be configured to automatically display on your screen if your computer has been inactive for a certain period of time. Originally used to prevent burn-in of monitors, screen savers are now generally used for entertainment or to password-protect users' desktops. You can access the screen saver options in Windows 11 by clicking Lock Screen under Personalization and then clicking Screen Saver Settings. The Screen Saver Control Panel tab also gives you access to other power settings.

Start The Start section allows you to set the options for the Start menu, such as to show the most used apps or to show recently added apps, just to name a few.

Taskbar The Taskbar section allows you to change the behavior of the Taskbar, such as automatically hiding it or using smaller buttons, just to name a few.

Fonts The Fonts section allows you to review any font installed on the operating system. These fonts are used by applications that allow changing of fonts, such as a Microsoft Office document. In addition, you can uninstall fonts and install new fonts from the Microsoft Store.



Windows 10/11 has replaced many of the Control Panel applets with the Settings app. If you cannot find a setting from a prior operating system in the Control Panel Desktop app, you can find the setting in the new Settings app.

In Exercise 1.1, you will see how to change a screen saver.

EXERCISE 1.1**Changing a Screen Saver in Windows**

1. Right-click the Desktop.
 2. Choose Personalize from the context menu.
 3. Click Lock Screen on the left side.
 4. Scroll to the bottom and select Screen Saver Settings.
 5. Choose a screen saver from the drop-down menu.
 6. Click Preview to view how the screen saver will look.
 7. Click OK to exit the Settings window.
-

The Taskbar

The *Taskbar* (see Figure 1.14) is another standard component of the Windows interface. Note that although the colors and feel of the Desktop components, including the Taskbar, have changed throughout the operating system versions, the components themselves are the same. In versions prior to Windows 11, the Taskbar contains two major items: the Start menu and the notification area, previously called the *system tray* (systray). The Start menu in Windows 11 is now left-center of the Taskbar and is easily identifiable: it is a button that has the Windows logo. The system tray is located on the right side of the Taskbar and contains a clock and date by default, speaker for sound volume, network connection status and OneDrive status, but other Windows utilities (for example, screen savers or antivirus utilities) may put their icons there to indicate that they are running and to provide the user with a quick way to access their features. On the far left of the Taskbar, you will find an icon that changes between, weather, sports, or headlines. If you click on the area, a news slide-out panel will fill the left side of the screen.

FIGURE 1.14 The Windows 11 Taskbar

Windows also uses the middle-right area of the Taskbar. When you open a new window or program, it gets a button on the Taskbar with an icon that represents the window or program. To bring that window or program to the front (or to maximize it if it was minimized), click its button on the Taskbar. As the middle-right area of the Taskbar fills with buttons, the buttons become smaller so that they can all be displayed.

Windows 11 allow you to pin commonly used programs to the Taskbar. The icon will appear on the Taskbar, and when the program is running, a short line will appear under the icon. You can pin any running task by right-clicking the icon in the Taskbar and selecting Pin To Taskbar. You can just as easily remove pinned icons by right-clicking them and selecting Unpin From Taskbar.

You can make the Taskbar automatically hide itself when it isn't being used (thus freeing that space for use by the Desktop or other windows). You can also make the Taskbar look like previous versions of Windows by selecting the alignment. In Exercise 1.2, we will show you how to do this.

EXERCISE 1.2

Modifying the Taskbar

1. Right-click the Taskbar.
2. Choose Taskbar Settings to bring up the Settings app for the Personalization options.
3. Click Taskbar Behaviors.
4. Select the option Automatically Hide The Taskbar. The Taskbar will automatically hide when the mode is turned on.
5. Select Taskbar Alignment and change it from Center to Left.
6. Move the mouse pointer to the bottom of the screen. The Taskbar will pop up and be available for normal use.


The Start Menu

Back when Microsoft officially introduced Windows 95, it bought the rights to use the Rolling Stones' song "Start Me Up" in its advertisements and at the introduction party. Microsoft chose that particular song because the Start menu was the central point of focus in the new Windows interface, as it was in all subsequent versions.

To display the Start menu, you can press the Windows key on your keyboard at any time. You can also click the Windows logo button in the Taskbar in Windows 11. You'll see a Start menu similar to the one shown in Figure 1.15 for Windows 11.



When Windows 8 was released, the Start menu was replaced with a Start screen. There was considerable backlash from the user community over the removal of the Start menu in Windows 8. This major change in the Start menu was partially responsible for Windows 8's slow adoption. A number of third-party vendors developed alternatives and marketed them to fill this demand.

FIGURE 1.15 Windows 11 Start menu

From the Start menu, you can select any pinned application. If you can't find your application in the pinned area, you can select All from the upper right of the menu. This will allow you to scroll down alphabetically to find your application. You can also use the search box to locate the application. Once you locate the application, if it is frequently used it is best to pin it to the Start menu. To pin an icon, move the mouse pointer over the icon and right-click. The context menu will appear, allowing you to pin or uninstall an application. You can also select More to open another context menu that allows you to Pin To The Taskbar, select the app settings, Rate and Review, or Share.



To check which OS you are using, press the Windows key and R key at the same time, type **winver.exe**, and then press Enter. The About Windows Application dialog will display the version and build number of the operating system.

Let's explore the principal features of the Windows 11 Start menu.

The Run Command

It is possible to run commands and utilities from the Start menu search box or from the Run dialog. To access the Run dialog box in Windows 11, press the Windows key and the R key at the same time. To execute a particular program, type its name in the Open field. If you don't know the exact path, you can browse to find the file by clicking the Browse button. Once you have typed the executable name, click OK to run the program.



To open a command prompt, type **command** in the Run box, and then click OK. You might need to run this as Administrator if you want to change system settings. From the Start menu, type **cmd** or **command** in the search field and then press Ctrl+Shift+Enter to run as Administrator.

Applications can easily be started from the Run window. You often will find it faster to open programs this way than to search for their icons in the Start menu maze. In Exercise 1.3, you will see how to start a program from the Run window.

EXERCISE 1.3

Starting a Program from the Run Window

1. Press the Windows key and the R key at the same time.
2. In the Open field, type **notepad**.
3. Click OK. Notepad will open in a new window.

If the program that you want to start has been run from the Run window before, you can find it on the Open field's drop-down list. Click the down arrow to display the list, and then select the program that you want by clicking its name and then OK.

Shut Down Command

Windows operating systems are very complex. At any one time, many files are open in memory. If you accidentally hit the power switch and turn off the computer while these files are open, there is a good chance that they will be corrupted. For this reason, Microsoft has added the Shut Down command under the Start menu. The command appears as an icon of a power button without a label. When you select this option, Windows presents you with several choices. The submenu will display Lock, Sleep, Shut Down, and Restart.

If the operating system requires security updates to install, there will be an orange dot in the lower-right corner of the power button icon. Additionally, the submenu will display Update And Shut Down and Update And Restart. If there is an optional feature update, the dot will be blue to grab your attention.

Recommended

The Recommended section of the Start menu displays recommended applications. It also shows frequently used applications based on your usage patterns.

Icons

Icons are shortcuts that allow a user to open a program or a utility without knowing where that program is located or how it needs to be configured. Icons consist of the following major elements:

- Icon label
- Icon graphic
- Program location or path

The label and graphic of the icon typically tell the user the name of the program and give a visual hint about what that program does. The icon for the Notepad program, for instance, is labeled Notepad and its graphic is a notepad. By right-clicking an icon once, you make it the active icon and a drop-down menu appears. One of the selections is Properties. Clicking Properties brings up the icon's attributes (see Figure 1.16), and it is the only way to see exactly which program an icon is configured to start and where the program's executable is located. You can also specify whether to run the program in a normal window or maximized or minimized.

Additional functionality has been added to an icon's properties to allow for backward compatibility with older versions of Windows (known as *compatibility mode*). To configure this, click the Compatibility tab and specify the version of Windows for which you want to configure compatibility. Note that you cannot configure compatibility if the program is part of the version of Windows that you are using. Figure 1.17 shows the compatibility settings available for an older program.

This feature is helpful if you own programs that used to work in older versions of Windows but no longer run under the current Windows version. In addition, you can specify different display settings that might be required by older programs.


Standard Desktop Icons

In addition to the options in your Start menu, a number of icons are placed directly on the Desktop in Windows. The Recycle Bin icon is one of these icons. In the latest version of Windows 11 (24H2), the Microsoft Edge icon can also be found on the Desktop. In older versions of Windows, the Computer icon could also be found on the desktop.


The Recycle Bin All files, directories, and programs in Windows are represented by icons. These icons are generally referred to as *objects*. When you want to remove

an object from Windows, you do so by deleting it. Deleting doesn't just remove the object, though; it also removes the ability of the system to access the information or application that the object represents. Therefore, Windows includes a special folder where all deleted files are placed: the Recycle Bin. The Recycle Bin holds the files until it is emptied or until you fill it. It gives users the opportunity to recover files that they delete accidentally. By right-clicking the Recycle Bin icon, you can see how much disk space is allocated. Some larger files that cannot fit in the Recycle Bin will be erased after a warning.

FIGURE 1.16 The Properties window of an application



You can retrieve a file that you have deleted by opening the Recycle Bin and then dragging the file from the Recycle Bin to where you want to restore it. Alternatively, you can right-click a file and select Restore. The file will be restored to the location from which it was deleted.

FIGURE 1.17 The Compatibility settings possible with an older program

In versions of Windows that this exam tests you on, the “deleted” files are stored in a folder called `\$Recycle.Bin`. This was not the case in all previous versions of Windows.

To erase files permanently, you need to empty the Recycle Bin, thereby deleting any items in it and freeing the hard drive space they took up. If you want to delete only specific files, you can select them in the Recycle Bin, right-click, and choose Delete. You can also permanently erase files (bypassing the Recycle Bin) by holding down the Shift key as you delete them (by dragging the file and dropping it in the Recycle Bin, pressing

the Del key, or clicking Delete on the file's context menu). If the Recycle Bin has files in it, its icon looks like a full trash can; when there are no files in it, it looks like an empty trash can.

Microsoft Edge If the eyes are the windows to the world, then Microsoft Edge is the window to the Internet (no pun intended). Because the Internet is such an integral part of a user experience on an operating system, it is placed on the desktop by default. Microsoft Edge is Microsoft's preferred browser of choice and it's based off the Chromium project.

This PC Icon If you double-click the This PC icon, it displays a list of all the disk drives installed in your computer. In addition to displaying disk drives, it displays a list of other devices attached to the computer, such as scanners, cameras, and mobile devices. The disk devices are sorted into categories such as Hard Disk Drives, Devices with Removable Storage, Scanners And Cameras, and so on.

You can delve deeper into each disk drive or device by double-clicking its icon. The contents are displayed in the same window.

In addition to allowing you access to your computer's files, the This PC icon on the Desktop lets you view your machine's configuration and hardware, also called the About page.



The This PC icon is not displayed on the Desktop by default. In Windows 11, you can add this icon by navigating to Settings > Personalization > Themes and clicking Desktop Icon Settings. You can then select the icons you want to display on the Desktop.


What's in a Window?

We have now looked at the nature of the Desktop, the Taskbar, the Start menu, and icons. Each of these items was created for the primary purpose of making access to user applications easier. These applications are, in turn, used and managed through the use of *windows*—the rectangular application environments for which the Windows family of operating systems is named. We will now examine how windows work and what they are made of.

A *program window* is a rectangular area created on the screen when an application is opened within Windows. This window can have a number of different forms, but most windows include at least a few basic elements.

Elements of a Window

Several basic elements are present in a standard window. Figure 1.18 shows the control box, title bar, Minimize/Maximize button, Close button, and resizable border in the text editor Notepad (`notepad.exe`) and Calculator (`calc.exe`) that has all the basic window elements—and little else.

FIGURE 1.18 The basic elements of a window, as seen in Notepad and Calculator

The basic window elements are as follows:

Control Box Located in the upper-left corner of the window, the control box is used to control the state of the application. It can be used to maximize, minimize, and close the application. Clicking it once brings into view a selection menu. Double-clicking it closes the window and shuts down the application. Some actions might differ from application to application, such as a double-click might maximize a window or the application.

Minimize and Maximize/Restore Buttons Used to change the state of the window on the Desktop. They are discussed in the section “States of a Window” later in this chapter.

Close Button Used to easily end a program and return any resources that it was using to the system. It essentially does the same thing as double-clicking the control box but with one fewer click.

Title Bar The area between the control box and the Minimize button. It states the name of the program and, in some cases, gives information about the particular document being accessed by that program. The color of the title bar indicates whether the window is the active window. Clicking and holding on it is an easy way to move the window on the screen.

Menu Bar Used to present useful commands within an application in an easily accessible format. Clicking one of the menu choices displays a list of related options from which you may choose.


Active Window The window that is currently being used. It has two attributes. First, any keystrokes that are entered are directed into the active window by default. Second, any other windows that overlap the active window are pushed behind it.

Border A thin line that surrounds the window in its restored down state and allows it to be resized.

Not every element is found in every window, because application programmers can choose to eliminate or modify each item. Still, in most cases, they will be consistent, with the rest of the window filled in with menus, toolbars, a workspace, or other application-specific elements. For instance, Microsoft Word, the program with which this book was written, adds a Ribbon control. It also has a menu bar, a number of optional toolbars, scroll bars at the right and bottom of the window, and a status bar at the very bottom. Application windows can become quite cluttered.

Notepad and Calculator are very simple Windows programs. Notepad has a single menu bar and the basic elements shown in Figure 1.18. Calculator has a simple menu that you can open by clicking the 3 lines in the upper-left of the application. Figure 1.19 shows a Microsoft Word window. Both Word and Notepad are used to create and edit documents, but Word is far more configurable and powerful and therefore has many more optional components available within its window.

FIGURE 1.19 A window with many more components, as seen in Microsoft Word



A quick way to access Help is to press the F1 key. In operating systems before Windows 10, a help and support application would launch. In Windows 10/11, the F1 key launches the web browser.

States of a Window

There is more to the Windows interface than the specific parts of a window. Windows also are movable, stackable, and resizable, and they can be hidden behind other windows (often unintentionally).

When an application window has been launched, it exists in one of three states:

Maximized A maximized window takes up all the available space on the screen. When it is in front of other programs, it is the only thing visible—even the Desktop is hidden. It takes up the entire space of the Desktop, and the middle button in the upper-right corner displays two rectangles rather than one. The sides of the window no longer have borders. The window is flush with the edges of the screen. Maximizing a window provides the maximum workspace possible for that window's application, and the user can actively access the window. In general, maximized mode is the preferred window size for most word processing, graphics creation, and other types of user applications.

Restored A restored window can be used interactively, and it is identical in function to a maximized window, with the simple difference that it does not necessarily take up the entire screen. Restored windows can be very small, or they can take up as much space as maximized windows. Generally, how large the restored window becomes is the user's choice. Restored windows display a Maximize button (the middle button in the upper-right corner) with a single rectangle in it; this is used to maximize the window. Restored windows have a border.

Minimized Minimized program windows are represented by nothing but an icon and title on the Taskbar, and they are not usable until they have been either maximized or restored. The difference between a minimized program and a closed program is that a minimized program is out of the way, but it is still taking up resources and is therefore ready to use if you need it. It also leaves the content of the window in the same place where it was when you minimized it.

When one program is open and you need to open another (or maybe you need to stop playing a game because your boss has entered the room), you have two choices. First, you can close the program currently in use and simply choose to reopen it later. If you do this, however, the contents of the window (your current game, for example) will be lost, and you will have to start over. Once the program has been closed, you can move on to open the second program.

The second option is to minimize the active window. Minimizing the game window, for example, removes the open window from the screen and leaves the program open but displays nothing more than an icon and title on the Taskbar. Later, you can restore the window to its previous size and finish the game in progress.



Keep in mind that applications in the background are still running. Therefore, if you minimize your game, you might return to find that you've been eaten by whatever monster you were running from in the game. A program running while minimized can be a good thing, however, if you're running a useful utility, such as a long search or an Internet download.

File Management

File management is the process by which a computer stores data and retrieves it from storage. Although some of the file-management interfaces across Windows may have a different look and feel, the process of managing files is similar across the board.

Files and Folders

In order for a program to run, it must be able to read information off the hard disk and write information back to the hard disk. To be able to organize and access information—especially in larger new systems that may have thousands of files—it is necessary to have a structure and an ordering process.

Windows provides this process by allowing you to create *directories*, also known as *folders*, in which to organize files. Windows also regulates the way that files are named and the properties of files. The filename for each file created in Windows has to follow certain rules, and any program that accesses files through Windows must also comply with these rules:

- Each file has a filename of up to 255 characters.
- Certain characters, such as a question mark (?) and slash (\ or /), are reserved for other uses and cannot be used in the filename. Periods are used to separate the filename from the extension, and the backslash is used to separate the directories in a path from the filename.
- A filename extension (generally three or four characters) can be added to identify the file's type.
- Filenames are not case sensitive. (You can create files with names that use both upper- and lowercase letters, but to identify the file within the filesystem, it is not necessary to adhere to the capitalization in the filename.) Thus, you cannot have a file named `working.txt` and another called `WORKING.TXT` in the same folder. To Windows, these filenames are identical, and you can't have two files with the same filename in the same folder. We'll get into more detail on this topic a little later.

In Windows 3.x and DOS, filenames were limited to eight characters and a three-character extension, separated by a period—known as the 8.3 file-naming convention. Windows 95 introduced long filenames, which allowed the 255-character filename convention.

The Windows filesystem is arranged like a filing cabinet. In a filing cabinet, paper is placed into folders, which are internal dividers, which are in a drawer of the filing cabinet. In the Windows filesystem, individual files are placed in subdirectories that are inside directories, which are stored on different disks or different partitions.

Windows also protects against duplicate filenames; no two files on the system can have exactly the same name and path. A *path* indicates the location of the file on the disk; it is composed of the letter of the logical drive on which the file is located and, if the file is located in a folder or subfolder, the names of those directories. For instance, if a file named `pagefile.sys` is located in the root of the C: drive—meaning it is not within a folder—the path to the

file is `C:\pagefile.sys`. As another example, if a file called `notepad.exe` is located under Windows under the root of `C:`, then the path to this file is `C:\Windows\notepad.exe`.



The *root folder* of any drive is the place where the hierarchy of folders for that drive begins. On a `C:` drive, for instance, `C:\` is the root folder.

Common filename extensions that you may encounter include `.exe` (for executable files, aka applications), `.dll` (for dynamic linked library files), `.sys` (for system files), `.log` (for log files), `.drv` (for driver files), and `.txt` (for text files). Note that DLL files contain additional functions and commands that applications can use and share. In addition, specific filename extensions are used for the documents created with each application. For example, the filenames for documents created in Microsoft Word have a `.doc` or `.docx` extension. You'll also encounter extensions such as `.mpg` (for video files), `.mp3` (for music files), `.png` and `.tif` (for graphics files), `.htm` and `.html` (for web pages), and so on. Being familiar with different filename extensions is helpful in working with the Windows filesystem.

Capabilities of File Explorer

Although it is technically possible to use the command-line utilities provided within the command prompt to manage your files, this generally is not the most efficient way to accomplish most tasks. The ability to use drag-and-drop techniques and other graphical tools to manage the filesystem makes the process far simpler. The *File Explorer* is a utility that allows you to accomplish a number of important file-related tasks from a single graphical interface.

Here are some of the tasks you can accomplish using File Explorer:

- View files and directories
- Open programs or data files
- Create directories and files
- Copy objects to other locations
- Move objects to other locations
- Delete or rename objects
- Change file attributes

You can access many of these functions by right-clicking a file or folder and selecting the appropriate option, such as Copy or Delete, from the context menu.



The File Explorer utility used in Windows 8/8.1 and Windows 10/11 was previously named Windows Explorer. Although the functionality remains almost identical since the introduction of Windows 95, the rebranding was just one of many changes when Windows 8 was released.

Navigating and Using File Explorer

Using File Explorer is simple. A few basic instructions are all you need to start working with it. First, the Explorer interface has a number of parts, each of which serves a specific purpose. The top area of File Explorer is dominated by a set of menus and toolbars that give you easy access to common commands. The main section of the window is divided into two panes: the left pane displays the drives and folders available, which is called the navigation pane, and the right pane displays the contents of the currently selected drive or folder, which is called the results pane. In recent versions of Windows, the Navigation pane is turned off. The following list describes some common actions in File Explorer:

Expanding a Folder You can click a folder in the left pane to expand it (show its subfolders in the left pane) and display its contents in the right pane. Clicking the right-pointing arrow (>) to the left of a folder expands the folder without changing the display in the right pane, as shown in Figure 1.20.

Collapsing a Folder Clicking the down arrow next to a folder collapses it.

Selecting a File If you click the file in the right pane, Windows highlights the file by marking it with a darker color.

Selecting Multiple Files The Ctrl or Shift key allow you to select multiple files at once. Holding down Ctrl while clicking individual files selects each new file while leaving the currently selected file(s) selected as well. Holding down Shift while selecting two files selects both of them and all the files in between.

Opening a File Double-clicking a file in the right pane opens the program if the file is an application; if it is a data file, it will open using the application for which the file-name extension is configured.

Changing the View Type Windows has several different view types: Extra Large Icons, Medium Icons, Large Icons, Small Icons, List, Details, Content, and Tiles. The choices vary a bit between the Windows versions. You can move between these views by clicking the View menu and selecting the view that you prefer. In Windows 10, the View menu is a ribbon that will drop down when clicked. The Navigation pane can be turned on from the View tab.

Finding Specific Files You access this option by using the Search button or bar. You can search for files based on their name, file size, file type, and other attributes.



When you're searching, you can also use wildcards. *Wildcards* are characters that act as placeholders for a character or set of characters, allowing, for instance, a search for all files with a .txt filename extension. To perform such a search, you'd type an asterisk (*) as a stand-in for the file-name: *.txt. An asterisk takes the place of any number of characters in a search. A question mark (?) takes the place of a single number or letter. For example, autoex??.bat would return the file autoexec.bat as part of its results.


Creating New Objects To create a new file, folder, or other object, navigate to the location where you want to create the object, and then right-click in the right pane (without selecting a file or folder). In the menu that appears, select **New** and then choose the object that you want to create.

Deleting Objects Select the object and press the **Del** key on the keyboard, or right-click the object and select **Delete** from the menu that appears.



The simplicity of deleting in Windows makes it likely that you or one of the people you support will delete or misplace a file or a number of files that are still needed. In such a case, the Recycle Bin (mentioned earlier) is a lifesaver.

FIGURE 1.20 Windows 11 File Explorer



Besides simplifying most file-management commands as shown here, File Explorer allows you to complete a number of disk-management tasks easily. For example, you can format and label removable media, which is discussed further in Volume 2 Chapter 3, “Windows Administration.”

Preparing for the Exam

Future chapters will delve further into operating systems and the tools, utilities, and features available with each. There is also additional coverage, as applicable, in the chapters on troubleshooting. For purposes of exam study, Table 1.3 offers a complete list of the features for each of the Windows operating systems that you need to know for the exam. It also details which chapter(s) in this book has more coverage of that particular topic.

TABLE 1.3 Windows features

Feature	Purpose	More information
BitLocker	Encrypts drives; available in each OS but not in every edition	Volume 2 Chapter 6
Domain access vs. workgroup	Shared security for computers and users	Volume 2 Chapter 3
Desktop styles/user interface	Customization of the desktop and the user interface	Volume 2 Chapter 2
Remote Desktop Protocol (RDP)	Allows users and administrators to connect remotely to obtain a desktop session	Volume 2 Chapter 8
Group Policy	A mechanism inside of Active Directory that allows for the management of user and computers	Volume 2 Chapters 2, 5

Summary

In this chapter, you learned about the basic operating systems, application installation, and the Windows 10/11 features. Additionally, we covered the basics of the Windows structure and window management. Because Windows is a graphical system, the key to success in learning to use it is to explore the system to find out what it can do. You will then be better prepared to decipher later what a user has done.

First, we explored the various operating systems you may encounter, along with their characteristics, such as their life cycles, categories, and minimum requirements.

Next, we discussed applications as they apply to the operating system and underlying architecture, as well as the various factors that should be evaluated before installing the application.

Finally, we introduced Windows 10/11 and its various editions. We then covered some basic Windows management concepts, including file management, as well as the folder structure. We also discussed using approved hardware and updating Windows.

With the basic knowledge gained in this chapter, you are now ready to learn how to interact with the most commonly used tools, the subject of the following chapter.

Exam Essentials

Know the RAM limitations between operating system architectures. 32-bit operating systems have a RAM limitation of 4 GB, whereas 64-bit operating systems can address over 4 GB of RAM.

Understand an operating system's life cycle. When development is started on an operating system, it is considered alpha or beta code. Release candidates precede the release to market (RTM) of an operating system. The end of mainstream support generally means only security updates will continue to be supported during the extended support period.

Know the various operating system types and the common operating systems. The four basic types of operating systems are server, workstation, mobile, and cloud-based. Windows, macOS, and Linux can be used as server, workstation, and mobile operating systems, whereas Google Chrome is primarily a cloud-based operating system.

Know the various editions of Windows 10/11 and their limitations. The four most common Windows 10/11 editions are Windows 10/11 Home, Pro, Pro for Workstations, and Enterprise. Windows 10/11 Home does not allow joining of domains and has many premium features. Windows 10 Pro is the most flexible edition because it allows for domain joining. Windows 10/11 Pro for Workstations allows for up to 4 CPUs and 6 TB of RAM. Windows 10/11 Enterprise is a volume licensed edition of Windows 10/11 that offers premium features.

Know the various requirements for applications. The system requirements for an application can be architecture, such as 32-bit versus 64-bit, dedicated graphics cards versus integrated, VRAM and RAM requirements, CPU requirements, the use of external licensing dongles or hardware tokens, storage requirements, and application compatibility.

Know the various cloud applications you will commonly support. You should be familiar with Exchange Online and Outlook terminology and basic setup. You should also know what each of the Microsoft 365 applications are used for, such as Word, Excel, PowerPoint, and Teams, as well as their Google counterparts.

Understand how to manage files in Windows. Nearly all file management—including moving, copying, renaming, and deleting files and changing file attributes—is accomplished through Windows File Explorer.

Know where files are located. The various versions of Windows that you need to know for this exam store files in multiple locations. You should be able to identify the location of those files mentioned in this chapter and be able to identify items such as where the Recycle Bin files are on each Windows operating system.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which of the following can you type in the Start menu in Windows to open a command prompt? (Choose two.)
 - A. run
 - B. cmd
 - C. command
 - D. open
2. Which part of the operating system can be described as extremely specific software written for the purpose of instructing the OS on how to access a piece of hardware?
 - A. Source code
 - B. Application
 - C. Kernel
 - D. Driver
3. Which application will sync your files to the cloud?
 - A. Excel
 - B. OneDrive
 - C. Exchange
 - D. Outlook
4. Which of the following installation methods would you use if you wanted to install a program on a virtual machine? (Choose the best method.)
 - A. ISO
 - B. USB
 - C. Optical disk
 - D. ZIP files
5. The Windows File Explorer program can be used to do which of the following? (Choose two.)
 - A. Browse the Internet
 - B. Copy and move files
 - C. Change file attributes
 - D. Create backup jobs
6. What is the maximum allowable length for the name of a file in Windows?
 - A. 8 characters
 - B. 32 characters

- C. 64 characters
 - D. 255 characters
7. Which of the following is a program that runs on top of the OS and allows the user to issue commands through a set of menus or some other graphical interface?
- A. Taskbar
 - B. Shell
 - C. Desktop
 - D. Source
8. If a program doesn't have a shortcut on the Desktop or in the All Apps submenu, you can start it by which of the following methods? (Choose the best answer.)
- A. Using the Shut Down command
 - B. Typing the program name in the Start box
 - C. Using the run command and typing in the name of the program
 - D. Typing **cmd** in the Start box followed by the program name
9. Which operating system feature offers the ability for a single application to have multiple requests into the processor at one time?
- A. Multiuser mode
 - B. Dystopia
 - C. Preemption
 - D. Multithreading
10. In Windows, a deleted file can be retrieved using which of the following?
- A. My Computer icon
 - B. Recycle Bin
 - C. Control Panel
 - D. Settings app
11. To turn off a Windows 11 operating system gracefully, you should do which of the following?
- A. Run the turnoff command at a command prompt.
 - B. Turn off the switch and unplug the machine.
 - C. Press Ctrl+Alt+Del.
 - D. Select Start, then click the power icon, and then choose Shut Down.
12. Which is required for a 64-bit installation of Windows 11?
- A. TPM 2.0
 - B. 32 GB Storage
 - C. Microsoft DirectX 9 graphics device
 - D. 4 GB RAM

13. What is the minimum amount of free hard drive space required for the installation of Windows 11 64-bit?
 - A. 1.5 GB
 - B. 15 GB
 - C. 32 GB
 - D. 64 GB
14. What is the minimum required memory for a 64-bit installation of Windows 11?
 - A. 512 MB
 - B. 1 GB
 - C. 2 GB
 - D. 4 GB
15. In Windows, a quick way to access Help is to press which keyboard key?
 - A. F12
 - B. The Windows key on the keyboard
 - C. F1
 - D. Alt
16. Which key(s) will open a run box?
 - A. Windows key+R
 - B. Ctrl+R
 - C. Windows key
 - D. F1
17. Which of the following is located on the rightmost portion of the Taskbar?
 - A. Start menu
 - B. Pinned apps
 - C. Notification area
 - D. Shutdown options
18. In addition to right-clicking the Desktop, how else can you access the Display Properties settings?
 - A. By clicking the Settings gear in the Start menu
 - B. By clicking the System icon under Control Panel
 - C. By pressing Ctrl+Alt+Esc
 - D. By pressing Ctrl+Alt+Tab

19. Which of the following allows you to navigate file/folder structures in Windows?
- A. Start menu
 - B. File Explorer
 - C. KDE
 - D. GNOME
20. What is the minimum required processor speed for a 64-bit installation of Windows 11?
- A. 2.2 GHz
 - B. 3.3 GHz
 - C. 1.0 GHz
 - D. 2.3 GHz

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

Your organization is planning to upgrade its operating system to Windows 11. Currently all of the laptops used in your organization use Windows 10 Pro. The organization eventually wants to roll out the Windows 11 feature of BranchCache. How should you proceed to accommodate the upgrade and the future feature needs?



Chapter 2

Windows Configuration

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **1.0 Operating Systems**
- ✓ **1.1 Explain common operating system (OS) types and their purposes.**
 - Various filesystem types
 - New Technology File System (NTFS)
 - Resilient File System (ReFS)
 - File Allocation Table 32 (FAT32)
 - Fourth extended filesystem (ext4)
 - Extended filesystem (XFS)
 - Apple File System (APFS)
 - Extensible File Allocation Table (exFAT)
- ✓ **1.2 Given a scenario, perform OS installations and upgrades in a diverse environment.**
 - Partitioning
 - GUID [globally unique identifier] Partition Table (GPT)
 - Master boot record (MBR)
 - Drive format
- ✓ **1.4 Given a scenario, use Microsoft Windows operating system features and tools.**
 - Task Manager
 - Services
 - Startup
 - Performance



- Processes
- Users
- Microsoft Management Console (MMC) snap-in
 - Event Viewer (eventvwr.msc)
 - Disk Management (diskmgmt.msc)
 - Task Scheduler (taskschd.msc)
 - Device Manager (devmgmt.msc)
 - Certificate Manager (certmgr.msc)
 - Local Users and Groups (lusrmgr.msc)
 - Performance Monitor (perfmon.msc)
 - Group Policy Editor (gpedit.msc)
- Additional tools
 - System Information (msinfo32.exe)
 - Resource Monitor (resmon.exe)
 - System Configuration (msconfig.exe)
 - Disk Cleanup (cleanmgr.exe)
 - Disk Defragment (dfrgui.exe)
 - Registry Editor (regedit.exe)

✓ **1.6 Given a scenario, configure Microsoft Windows settings.**

- Internet Options
- Devices and Printers
- Programs and Features
- Network and Sharing Center
- System
- Windows Defender Firewall
- Mail
- Sound
- User Accounts
- Device Manager



- Indexing Options
- Administrative Tools
- File Explorer Options
 - View hidden files
 - Hide extensions
 - General options
 - View options
- Power Options
 - Hibernate
 - Power plans
 - Sleep/suspend
 - Standby
 - Choose what closing the lid does
 - Turn on fast startup
 - USB selective suspend
- Ease of Access
- Time and Language
- Update and Security
- Personalization
- Apps
- Privacy
- System
- Devices
- Network and Internet
- Gaming
- Accounts



The previous chapter introduced the basic components of the Windows operating systems. This chapter builds on that and focuses on the configuration of the Windows 11 operating system.

Interacting with Operating Systems

In the following sections, we will look at the Microsoft GUI from the ground up. In Volume 2 Chapter 1, “Operating System Basics,” we took a detailed look at its key components, and we will build on that with an exploration of basic tasks for Windows 11. Although the exam objectives are focused on Windows 11, most of these configuration concepts can be applied to Windows 10, and even to old Windows operating systems in many cases.

Microsoft has included a number of tools with each iteration of Windows to simplify system configuration. Although some tools have specific purposes and are used only on rare occasions, you will come to rely on a number of tools and access them on a regular basis. It is this latter set that we will examine in the following sections.



Exam Objective 1.4, “Given a scenario, use Microsoft Windows operating system features and tools,” requires you to know the elements of the Task Manager, various Microsoft Management Console (MMC) snap-ins, and additional tools. Be sure to know how to use the Task Manager, and know what each of the MMC snap-ins and additional tools provide.

Task Manager


Task Manager lets you shut down nonresponsive applications selectively in all Windows versions. In current versions of Windows, it can do so much more, allowing you to see which processes and applications are using the most system resources, view network usage, see connected users, and so on. To display Task Manager, press Ctrl+Alt+Delete and click the Task Manager button or option. You can also right-click an empty spot in the Taskbar and choose Task Manager from the context menu.



To get to Task Manager directly in any of the Windows versions that include it, you can press Ctrl+Shift+Esc.

Depending on the Windows version, Task Manager has various tabs or views. Figure 2.1 shows the common default display in Windows 11, but other versions vary from the seven tabs shown here: Processes, Performance, App History, Startup Apps, Users, Details, and Services. In the following figure the side menu is expanded by clicking the three lines in the upper-left of the Task Manager. By default the menu is collapsed and only the icons are displayed.

FIGURE 2.1 The default Task Manager in Windows




Let's look at these tabs in more detail:

Processes The Processes tab lets you see which tasks are open on the machine. You also see the status of each task, which can be either Running or Not Responding. If a task/application has stopped responding (that is, it's hung), you can right-click the task in the list and click End Task. Doing so closes the program, and you can try to open it again. Often, although certainly not always, if an application hangs, you have to reboot the computer to prevent the same thing from happening again shortly after you restart the application. You can also use the Processes tab to switch to a different task or create new tasks. In Windows 11, many of the elements that appear in Processes are copied between the Processes and Details tabs; an App History tab was added that, as the name implies, displays usage settings for applications and the currently logged-in user account.

Details In Windows 11 the Details tab, shown in Figure 2.2, gives you more detail than the Processes tab. The Details tab can display a multitude of metrics and details, such as

Program ID (PID), Status, User Name, and UAC Virtualization Status, and these are just the defaults. You can select the columns you want to see by right-clicking on the header, clicking Select Columns, then choosing the columns you want to inspect, as shown in Figure 2.3.

FIGURE 2.2 The Details tab of Task Manager



To end a process, right-click it in the list and click End Task. Be careful with this choice, because ending some processes can cause Windows to shut down. If you don't know what a particular process does, you can look for it in any search engine and find a number of sites that explain it.

You can also change the priority of a process in Task Manager's Details tab by right-clicking the name of the process and choosing Set Priority. The six priorities, from lowest to highest, are as follows:

Low For applications that need to complete sometime but that you don't want interfering with other applications. On a numerical scale from 0 to 31, this equates to a base priority of 4.


Below Normal For applications that don't need to drop all the way down to Low. This equates to a base priority of 6.

Above Normal For applications that don't need to boost all the way to High. This equates to a base priority of 10.

High For applications that must complete soon, when you don't want other applications to interfere with the applications' performance. This equates to a base priority of 13.

Realtime For applications that must have the processor's attention to handle time-critical tasks. Applications can be run at this priority only by a member of the Administrators group. This equates to a base priority of 24.

FIGURE 2.3 The various columns for the Details tab




If you decide to change the priority of an application, you'll be warned that doing so may make it unstable. You can generally ignore this option when changing the priority to Low, Below Normal, Above Normal, or High, but you should heed this warning when changing applications to the Realtime priority. Realtime means that the processor gives precedence to this process over all others—over security processes, over spooling, over everything—and this is sure to make the system unstable.

Task Manager changes the priority only for that instance of the running application. The next time the process is started, priorities revert back to that of the base (typically, Normal).

Services The Services tab lists the name of each running service as well as the process ID associated with it and its description, status, and group. A button labeled Open Services appears on the bottom of this screen. Clicking it will open the *Microsoft Management Console (MMC)* snap-in for Services, where you can configure each service. Within Task Manager, right-clicking a service will open a context menu listing six choices: Start, Stop, Restart, Open Services (which opens the Service MMC), Search Online, and Go To Details (which takes you to the Processes or Details tab).

Performance The Performance tab displays a variety of information, including overall CPU usage percentage, a graphical display of CPU usage history, the number of processes, and a graphical display of physical memory. This tab also provides you with additional memory-related information, such as physical and kernel memory usage as well as the total number of handles, threads, and processes. Total, limit, and peak commit-charge information also appears, as shown in Figure 2.4.

FIGURE 2.4 The Performance tab of Task Manager



Some of the items are beyond the scope of this book, but it's good to know that you can use the Performance tab to keep track of system performance. Note that the number of processes, CPU usage percentage, and commit-charge information always appear at the bottom of the Task Manager window, regardless of which tab you have currently selected.

Users The Users tab provides you with information about the users logged into the local machine. You'll see the username, status, CPU usage, memory usage, disk usage, and network usage. You can right-click the name of any connected user to perform a variety of functions, including viewing active processes, sending the user a message, signing off the user, switching accounts, and initiating a remote-control session to the user's machine, if they are connected remotely.

Startup The Startup tab lists the names of services configured to begin at startup as well as the publisher, status, and startup impact. From this tab, you can select any service listed and choose to disable it. This functionality was provided by the `msconfig` command in prior versions of Windows, such as Windows 7.

Use Task Manager whenever the system seems bogged down by an unresponsive application. You will work with Task Manager in Exercise 2.1.

EXERCISE 2.1

Working with Task Manager

1. Select Start, type **notepad**, and then press Enter.
2. Right-click an empty portion of the Taskbar and select Task Manager.
3. Find Notepad on the Processes tab, right-click, and select Go To Details.
4. Right-click on `notepad.exe` on the Details tab that opened and select Set Priority, then Above Normal.
5. Confirm the change by clicking Change Priority in the dialog box that opens to confirm your change.
6. Right-click on the column header and choose Select Columns from the context menu.
7. Scroll through the list and select the Base Priority option, then click OK.
8. Observe the status of `notepad.exe`.
9. Right-click `notepad.exe` and select End Task.
10. Confirm the ending of the task by clicking End Process in the dialog box that displays.

Microsoft Management Console

Microsoft created the Microsoft Management Console (MMC) interface as a frontend in which you can run administrative and configuration tools. Many administrators don't even know that applications they use regularly run within an MMC. In the following sections, we will cover many of the different MMC snap-ins you will use in your day-to-day

administration and configuration of the operating system. You can start the MMC by pressing Windows key+R to open Run, typing **mmc.exe** in the box, and then clicking OK. Once the MMC is started, you can create a custom MMC, adding the snap-ins discussed next. Click File ➤ Add Or Remove Snap-ins, and then select which snap-ins to add. For more information about the MMC, visit <https://learn.microsoft.com/en-us/troubleshoot/windows-server/system-management-components/what-is-microsoft-management-console>.


Computer Management

Windows includes a snap-in to manage computer settings: the Computer Management console. The Computer Management console can manage more than just the installed hardware devices; it can manage all the services running on a computer in addition to Device Manager, which functions almost identically to the one that has existed since Windows 9x. It also contains Event Viewer, which shows any system errors and events as well as methods to configure the software components of all the computer's hardware.

To access the Computer Management console, right-click the Start menu and select Computer Management from the context menu. You can also press Windows key+R and type **compmgmt.msc** to start the console.

After you are in Computer Management, you will see all the tools available, as shown in Figure 2.5. This is one power-packed snap-in and includes the following system tools:

FIGURE 2.5 Computer Management



Task Scheduler Allows for the scheduling of tasks using a simple command, script, or external program. Many of the operating system's functions are tasks such as checking for updates.

Event Viewer A link to the tool that allows you to view application error logs, security audit records, and system errors. We cover Event Viewer in detail shortly.

Shared Folders Allows you to manage all your computer's shared folders. You can also view active remote sessions and open files from this console.

Local Users and Groups Allows you to create and manage local user and group accounts.

Performance Shows you how your system hardware is performing and alerts you if system performance goes under a threshold that you set.

Device Manager Allows you to manage hardware devices and perform tasks such as updating drivers, disabling devices, and uninstalling devices. Device Manager is the first place to visit if a device is physically installed but not working.

Storage Allows you to perform basic disk management on the operating system.

Services and Applications Allows you to control services in the operating system.

Event Viewer


Event Viewer (`eventvwr.msc`) is an MMC snap-in that shows a lot of detailed information about what is running on your operating system. You can start it in several different ways—for example, by clicking Start, typing **Event Viewer**, and selecting the Event Viewer app, or by right-clicking the Start button and selecting Event Viewer from the context menu. In addition, you can add it as a snap-in inside the MMC or press Windows key+R, type **eventvwr.msc**, and press Enter.

Event Viewer should be the first place you look when you are trying to solve a problem whose solution is not evident. The system and applications will often create an entry in Event Viewer that can be used to verify operation or diagnosis problems, as shown in Figure 2.6.

Table 2.1 highlights the three main event logs that you should be concerned with for the exam. Each feature in Windows has the ability to store its events in specific log files, in application and service logs.

TABLE 2.1 Event Viewer logs

Event Log	Description
Application	Events generated by applications installed on the operating system
Security	Events generated by the Security Reference Monitor in the Executive kernel
System	Events generated by the operating system


FIGURE 2.6 Event Viewer

Although you might think that all the security-related information is in the Security log, you're only half right. The Security log is used by the Security Reference Monitor inside the *Executive kernel*. It is responsible for reporting object audit attempts. Examples of object audit attempts include file access, group membership, and password changes.

Most of the useful security-related information will be in the application and system logs. Using these logs, you can see errors and warnings that will alert you to potential security-related problems. When you suspect an issue with the operating system or an application that interacts with the operating system, you should check these logs for clues. The event log won't tell you exactly what is wrong and how to fix it, but it will tell you if there is an issue that needs to be investigated.

Disk Management

The Disk Management (`diskmgmt.msc`) snap-in is used to view disk information, such as volumes configured on the physical disk and the *filesystems* that are formatted on the volume. Disk Management, shown in Figure 2.7, isn't used just to view information; you can also use it to partition volumes on a new or existing disk, format filesystems, and mount volumes to drive letters. These are just a few configuration tasks; we will cover Disk Management later in this chapter.

FIGURE 2.7 Disk Management

Task Scheduler

Task Scheduler (`taskschd.msc`) is accessible beneath the Computer Management snap-in or inside Windows Tools from the Start menu under All Apps. This snap-in will allow you to configure an application to run automatically or at any regular interval (see Figure 2.8). A number of terms are used to describe the options for configuring tasks: *action* (what the task actually does), *condition* (an optional requirement that must be met before a task runs), *setting* (any property that affects the behavior of a task), and *trigger* (the required condition for the task to run).


For example, you could configure a report to run automatically (action) every Tuesday (trigger) when the system has been idle for 10 minutes (condition), and only when requested (setting).

Device Manager

Device Manager (`devmgmt.msc`) is an indispensable tool for the management of peripherals and components attached to the computer. You can view all the devices in the system, as shown in Figure 2.9. Device Manager has been around since Windows 95, and it hasn't changed all that much. Because computer systems have evolved and all newer computer

systems are plug-and-play capable, this tool is usually only used when you are in doubt that a component is working or you're troubleshooting a problem, or when you want to check the driver. Device Manager was used in the past a lot more since most peripherals had to be manually configured.

FIGURE 2.8 Task Scheduler



Device Manager allows you to manually update the driver for a device, roll back a driver to a prior version, uninstall a device, and disable a device. Some peripherals, such as the network card, allow you to configure if you want the device to wake the computer. Some peripherals, such as a disk drive or USB device, can be configured to fall asleep if they are not used for a specified period of time.

There are a number of states that a device inside of Device Manager can be in. The following are the possible states:

Working In a working state the device will not have any symbol or icon overlaid and the device will function as normal.


Disabled In a disabled state the device will not function. The symbol or icon overlaid will be a down arrow in a circle.

Problem In a problem state the device can function, but there is an underlying issue. The symbol or icon overlaid on the device will be a yellow exclamation mark in either a circle or triangle.

Driver Issue In a missing driver issue state the device is functioning with a compatible driver and lacks a device-specific driver. The symbol or icon overlaid will be a question mark.

Manual In a manual state the device can function as normal. The symbol or icon overlayed will be in the form of a lowercase i. This signifies that the Automatic Settings option is not selected for the device.

FIGURE 2.9 Device Manager




Certificate Manager

Certificate Manager (`certmgr.msc`) is used to view and manage certificates used by the web browser and the operating system, as shown in Figure 2.10. Certificates are used for a number of reasons; the most common are encryption and digital signatures that provide trust.

Certificate Manager allows you to manage certificates for your user account, a service account, or the computer account. When you choose to manage your user account, the certificates managed can be used only by your account. So, the certificates are only relevant while you are logged in and using an application that requires a certificate from the certificate store. When you choose to manage certificates for a particular service account, the certificate is only relevant to that specific service, such as a virtual private network (VPN) service. When you choose to manage certificates for the computer account, these

certificates are relevant only for the operating system, even if someone is not logged in. This configuration mode is commonly used when configuring a certificate for the Internet Information Services (IIS) web server which is integrated into most editions of Windows Server.

FIGURE 2.10 Certificate Manager




Local Users and Groups

The Local Users and Groups (lusrmgr.msc) MMC snap-in allows for granular control over local user accounts and groups for the Windows operating system. You can access the Local Users and Groups MMC by right-clicking the Start menu and choosing Computer Management. You can also press the Windows key+R, type **lusrmgr.msc**, and press Enter. When the tool launches, if you click Users you will see several built-in user accounts, all of which are disabled (depicted with the down arrow), as shown in Figure 2.11. The only active account on a brand-new Windows 11 operating system is the first account setup.


The MMC extension is divided into the following two parts:

- The Users section allows for the viewing, creation, and modification of individual users. You can also access disabled users, such as the built-in Administrator account and Guest account.
- The Groups section allows for the viewing, creation, and modification of groups. The User Accounts Control Panel applet only allows assignment of the administrator group with a radio button–type dialog box. However, from this MMC extension, you can assign many different groups. Each group assigns more permissions to a user.

FIGURE 2.11 Local Users and Groups

Performance Monitor

Performance Monitor (perfmon.msc) varies a bit in different versions of Windows, but it has the same purpose throughout: to display performance counters, as shown in Figure 2.12. The tool collects counter information and then sends that information to a console or event log.

FIGURE 2.12 Performance Monitor

Performance Monitor's objects and counters are very specific; you can use Performance Monitor as a general troubleshooting tool as well as a security troubleshooting tool. For instance, you can see where resources are being used and where the activity is coming from. In Exercise 2.2, you see how to work with Performance Monitor.

EXERCISE 2.2

Working with Performance Monitor

1. Select Start, type **perfmon**, and then press Enter.

Make sure that you select Performance Monitor, as opposed to System Monitor or another default that may come to the forefront.

2. Choose the Performance Monitor section under Monitoring Tools.
3. Click the plus sign (+) or right-click in the graphical display area and select Add Counters.
4. Expand the Processor section, and then select the %ProcessorTime object.
5. Click Add and then click OK.
6. Open Windows File Explorer, click the C: drive, type * into the search box, and then press Enter.
7. Quickly change to Performance Monitor and watch the impact of this search on the processor.

This action is time consuming and therefore will help you notice the changes that take place in Performance Monitor.

8. Run the same operation again. This time, however, change your view within Performance Monitor to the histogram bar by clicking the button directly to the left of the plus sign.
 9. Run the same operation again, changing your view within Performance Monitor to Report.
 10. Exit Performance Monitor.
-


Group Policy Editor

The Group Policy Editor (`gpedit.msc`) tool allows you to edit the local Group Policy for the operating system. *Group Policy* is a mechanism that allows an administrator to set various settings to customize the operating system. These settings can restrict the operating


system in a number of ways, such as removing the Settings tab from a built-in application. Group Policy can also control aspects of security for the operating system. Group Policy also contains a mechanism to enforce these settings by reapplying the settings periodically in the event they are changed.

The Group Policy Editor divides all settings between computer settings and user settings. The computer settings affect the operating system and how it behaves. The user settings affect the user logged onto the operating system. If you open the Group Policy Editor by typing **gpedit.msc** in the Run dialog box and pressing Enter, the tool will open and display a tree of settings, as shown in Figure 2.13.

FIGURE 2.13 Group Policy Editor



There is another way to open the Group Policy Editor that allows you to edit the Group Policy for a specific user, administrators, or all non-administrators. You must first start the MMC tool by typing **mmc.exe** in the Run dialog box and then pressing Enter. Then, add the snap-in Group Policy Object Editor. After it is selected, another dialog box will pop up where you can browse to select a specific user or group or a special group of users, as shown in Figure 2.14.

FIGURE 2.14 Group Policy Editor browse dialog box

Additional Tools

When the Microsoft Management Console (MMC) was first introduced with Windows 2000, it was to be a single pane of glass for monitoring and configuration of the Windows operating system. Almost 25 years later, we still use a mixture of tools outside the MMC to monitor and configure Windows. Next, we discuss some of these additional utilities and tools.

System Configuration Tool

The System Configuration (`msconfig.exe`) tool allows you to configure how Windows 11 starts up, as well as launching additional tools. The tabs of the System Configuration tool differ a bit based on the Windows version you are running. The main tabs are General, Boot, Services, Startup, and Tools. Figure 2.15 shows the General tab for Windows 11. From here, you can configure the startup options.

Figure 2.16 shows the Boot tab for Windows 11. Note that from here, you can configure the next boot to be a safe boot, and you can turn on the boot information so that you can see drivers as they load—which is quite useful when a system keeps hanging during boot. Although the option is titled Safe Boot, this is classically referred to as booting into Safe mode.



FIGURE 2.15 System Configuration General tab**FIGURE 2.16** System Configuration Boot tab

Figure 2.17 shows the Services tab for Windows 11. On this tab, you can view the services installed on the system and their current status (running or stopped). You can also enable or disable services.

In Windows 7 and earlier, the Startup tab allowed you to configure applications that start up when any user logs in. In Windows 10/11, the Startup tab redirects you to the Startup tab in Task Manager, where these tasks can be performed, as shown in Figure 2.18.

FIGURE 2.17 System Configuration Services tab




FIGURE 2.18 System Configuration Startup tab and Task Manager Startup tab




Figure 2.19 shows the Tools tab for Windows 11. On this tab, you can launch a number of administrative tools to configure various Windows features.

Keep in mind that the tabs differ slightly based on the operating system version. We walked through the CompTIA objectives related to this tool in this discussion.

System Information Tool

The System Information (`msinfo32.exe`) tool displays a fairly thorough list of settings on the machine (see Figure 2.20). You cannot change any values here, but you can search, export, and save the output. Several command-line options can be used when starting `msinfo32`; Table 2.2 summarizes them.

FIGURE 2.19 System Configuration Tools tab




FIGURE 2.20 The Msinfo32 interface shows configuration values for the system.




TABLE 2.2 msinfo32 command-line options

Option	Function
/computer	Allows you to specify a remote computer on which to run the utility
/nfo	Creates a file and saves it with an .nfo extension
/report	Creates a file and saves it with a .txt extension



The use of command-line options will be covered in Volume 2 Chapter 3, “Windows Administration,” when you learn about installing and upgrading to Windows 11.

Resource Monitor

Resource Monitor (`resmon.exe`) is used to identify resource utilization of CPU, disk, network, and memory on Windows. The utility was originally introduced in Windows Vista and has been included with all new releases of Windows. The utility can be launched a few different ways. The first is by entering the command **resmon.exe** in the Run dialog box. Another way to open the Resource Monitor is from inside Task Manager. The upper-right of the Performance tab is three dots, if you click on them and then select the option Resource Monitor; clicking it will launch Resource Monitor in a new window. Figure 2.21 shows Resource Monitor and the four tabs for each resource.

The CPU tab allows you to identify the process with the highest amount of CPU utilization on the operating system. Each column can be sorted; it is as simple as clicking the CPU column to sort the column and find the process with the highest CPU utilization. When you do find the process, you can right-click the entry and choose to end the process, suspend the process, and even search the process online.


The Memory tab displays detailed memory usage of the processes running on the operating system, as shown in Figure 2.22. This tab can help you identify memory in use, memory reserved by hardware, and memory available by the operating system and processes.

The Disk tab helps you identify a process that is overusing the hard drive with a high amount of read requests, write requests, or overall usage. The Disk tab will also allow you to identify the I/O priority of processes and their response time. This tab is extremely useful when you suspect that a process is slowing down the system.

The Network tab displays all the processes that are currently utilizing the network. The processes can be sorted by send, received, and total bytes per second. The Network tab does a lot more than just displaying activity; it also shows the destination addresses for each

process. This is valuable information if you suspect name resolution is a problem with the remote application. Opening the TCP Connections drop-down, you can view the active TCP connections on the operating system, along with packet loss and latency. Normally this active view of network traffic can only be performed with a packet capture tool or another third-party tool. The Network tab is not just useful for information on outgoing applications, it can also display processes listening on TCP and UDP ports. The Resource Monitor tool will also display *firewall* status for the processes listening on the operating system.

FIGURE 2.21 Resource Monitor



Disk Cleanup

Storage space is finite in a computer, and it is inevitable that it will be consumed over time. It can be consumed by a number of files, such as the operating system, downloaded programs, temporary files, updates, deleted files, and of course your data files. The Disk Cleanup (`cleanmgr.exe`) tool can clean up operating system files to free up space without affecting your data files. Disk Cleanup, shown in Figure 2.23, can be launched by entering the command **cleanmgr.exe** in the Run dialog box, or by searching for Disk Cleanup on the Start menu.

FIGURE 2.22 Resource Monitor Memory tab





FIGURE 2.23 Disk Cleanup



Once Disk Cleanup is launched, you can select the various files you want to clean up on the system drive. The categories you can delete are Downloaded Program Files, Temporary Internet Files, DirectX Shader Cache, Delivery Optimization Files, Recycle Bin, Temporary Files, Thumbnails, Windows Update Files, Windows Defender Antivirus, and countless other files. For some categories, you can view the files that you can potentially delete; two such categories are Temporary Internet Files and Downloaded Program Files. However, there are other categories in which you can view the individual files. After selecting the categories you want to delete, you can purge them from the disk, thus freeing up space.



Objective 1.6, “Given a scenario, configure Microsoft Windows settings,” requires you to know various Control Panel settings and their controls.

Control Panel

Although, for the most part, Windows is functional from the time it is installed, Microsoft realized that if someone were going to use computers regularly, they would probably want to be able to customize their environment so that it would be better suited to their needs—or at least more fun to use. As a result, the Windows environment has a large number of utilities that are intended to give you control over the look and feel of the operating system.

This is, of course, an excellent idea. It is also a bit more freedom than some less-than-cautious users seem to be capable of handling. You will undoubtedly serve a number of customers who call you to restore their configuration after botched attempts at changing one setting or another.

More than likely, you will also have to reinstall Windows yourself a few times because of accidents that occur while you are studying or testing the system’s limits. This is actually a good thing, because no competent computer technician can say that they have never had to reinstall because of an error. You can’t really know how to fix Windows until you are experienced at breaking it. So, it is extremely important to experiment and find out what can be changed in the Windows environment, what results from those changes, and how to undo any unwanted results. To this end, we will examine the most common configuration utility in Windows: Control Panel, as shown in Table 2.3, which describes some popular Control Panel applets. Also, not all applets are available in all versions and many of them have gradually been moved by Microsoft to the Settings.

TABLE 2.3 Selected Windows Control Panel applets

Applet name	Function
Device Manager	Adds and configures new hardware.
Programs and Features	Changes, adds, or deletes software.
Administrative Tools	Performs administrative tasks on Windows 10. In Windows 11 it has been replaced with Windows Tools.

TABLE 2.3 Selected Windows Control Panel applets *(continued)*


Applet name	Function
File Explorer Options	Configures the look and feel of how folders are displayed in Windows File Explorer.
Internet Options	Sets a number of Internet connectivity options.
Network and Sharing Center	Sets options for connecting to other computers.
Power Options	Configures different power schemes to adjust power consumption.
Devices and Printers	Configures printer settings and print defaults.
System	Allows you to view and configure various system elements (discussed in more detail later in this chapter).
Windows Defender Firewall	Configures basic firewall exemptions.
Mail (Microsoft Outlook)	Configures the Outlook mail profile (only appears if Outlook is installed).
Sound	Configures audio.
User Accounts	Configures local accounts on the operating system.
Indexing Options	Configures the folders to index for search capabilities.
Ease of Access Center	Allows for accessibility options for users.

In the current version of Windows, when you first open Control Panel, it appears in Category view, as shown in Figure 2.24. Control Panel programs have been organized into various categories, and this view provides you with the categories from which you can choose. When you choose a category and pick a task, the appropriate Control Panel program opens. Or, you can select one of the Control Panel programs that is part of the category.

You can change this view to Small Icons or Large Icons in Windows 11, which displays all the Control Panel programs in a list, as in older versions of Windows. The items are organized by category and not by large icons in Windows by default. Therefore, we *strongly*

suggest that administrators change to the large icon view. To do so, select Large Icons in the View By drop-down box in the right corner of Control Panel. Throughout this chapter, when we refer to accessing Control Panel programs, we will assume that you have changed the view to the Large Icons view.

FIGURE 2.24 The Windows Control Panel in Category view




For a quick look at how the Control Panel programs work, in Exercise 2.3, you'll examine some of the settings in the Date and Time applet (`timedate.cpl`).



Filenames appearing in parentheses, such as `timedate.cpl` in the preceding paragraph, can be used to start the utility without needing to go through the Control Panel dialog boxes.

The Date and Time applet is used to configure the system time, date, and time zone settings, which can be important for files that require accurate time stamps (or for users who don't wear a watch). Because Date and Time is a simple program, it's a perfect example to use. Current versions of Windows have an Internet Time Settings tab, which enables you to synchronize time on the computer with an Internet time server (the options in Windows 11 are shown in Figure 2.25). By default, the Internet time server is set to `time.windows.com`.

FIGURE 2.25 Windows Date and Time/Internet Time settings**EXERCISE 2.3****Changing the Time Zone**

1. Click Start, type **Control Panel** or **timedate.cpl**, and then select the appropriate result from the search.
2. From Control Panel, double-click the Date and Time icon. (By default, the programs are listed alphabetically.)
3. Click the Change Time Zone button, and from the drop-down menu, select (GMT -03:30) Newfoundland.
4. Hop a plane to Newfoundland, secure in the knowledge that you will know what time it is once you get there.
5. If you skipped step 4, change the time zone back to where it should be before closing the window.




Microsoft is slowly moving away from Control Panel in favor of the Settings app. Many of the same settings appear in both Control Panel and the Settings app. Some settings can only be set in the Settings app, such as Windows Update settings, whereas others, such as Storage Spaces, can only be set in the classic Control Panel.

Regional and Language Options

You can configure regional settings through the Control Panel applet Region in Windows 11. Using this applet (`intl.cpl`), you can choose which format is used for numbers (see Figure 2.26), your geographic location, and the language to be used for non-Unicode programs. In Windows 11, the Settings app is also used to set the layout of the keyboard you are using as well as set the regional settings previously mentioned. We will explore the Settings app further on in this chapter.

FIGURE 2.26 Windows Region Control Panel applet




The ability to support so many languages is provided through the use of the *Unicode* standard. In Unicode, and the Unicode Character Set (UCS), each character has a 16-bit value. This allows the same character to be interpreted/represented by 65,536 different entities.

If you click Additional Settings, you can go beyond the date and time formats and also configure number and currency, as shown in Figure 2.27.

Internet Options Applet

The Internet Options applet (`inetcpl.cpl`) brings up the Internet Properties dialog box, as shown in Figure 2.28. The tabs include General, Security, Privacy, Content, Connections, Programs, and Advanced. Use this applet to configure the browser environment for any program running on Windows that makes use of Internet Explorer settings, such as user proxy settings or Internet security zone setting, just to name a few.

FIGURE 2.27 Windows Additional Settings Region Control Panel applet


It's important to note that in Windows 10, Internet Explorer exists alongside the Microsoft Edge browser. Microsoft has incorporated an Internet Explorer mode for Microsoft Edge; this will allow a smoother transition to Edge. The sunset date (EOL) for Internet Explorer was June 15, 2022, and Windows 11 does not include Internet Explorer. The settings in this Control Panel applet can also be set in the Settings app.

File Explorer Options Applet

The File Explorer Options applet will open to the General tab, as shown in Figure 2.29. Using this tab, you can change the default opening pane for File Explorer. The Home option, the default, displays quick access shortcuts, favorites, and recent files. The This PC option displays the common folders found under This PC, along with devices and drives.

In addition, you can control how folders open, such as opening folders in the same window or opening folders in their own new windows. You can change how items are opened, such as single-click or double-click. This tab also lets you change privacy settings, such as showing recently used file folders in the Home view. By default, Windows 11 will be helpful by showing these recently used files and folders, but you may want to shut that behavior off. After doing so, you should clear the File Explorer history by clicking the Clear button in the General tab of the Folder Options applet.


FIGURE 2.28 Internet Options Control Panel applet




The View tab in the File Explorer Options applet allows you to change how files and folders are viewed in File Explorer. There are a number of settings on this tab that will allow you to change the way File Explorer is viewed. The settings range from always showing menus to showing all folders in the Navigation pane. One of the first settings that is usually changed is Hide Extensions For Known File Types, because seeing the extensions is really handy.



In later versions of Windows 10/11, the ability to see extensions is the default. This was changed because many malicious files can hide their extensions in an attempt to be executed or opened.

FIGURE 2.29 File Explorer General Options

Some of the more important files that you will need to work on are hidden by default as a security precaution. To make certain folders or files visible, you need to change the display properties of Windows File Explorer, as shown in Figure 2.30. You learn how to do this in Exercise 2.4.

FIGURE 2.30 File Explorer View Options

EXERCISE 2.4**Showing Hidden Files and Folders**

1. Open Windows File Explorer on a Windows system.
2. Browse to the root of the C: drive and note the entries that are there.
3. If it is available, choose View, then Show, and select the Hidden Items option so that they are now visible.
4. Note the entries that now appear but that did not previously.
5. Open the File Explorer Options applet from Control Panel and select the View tab, if needed.
6. Deselect Hide Protected Operating System Files (Recommended).
7. Deselect Hide File Extensions For Known File Types.
8. Click OK.

You will now be able to see the Windows system files discussed in the following sections. For security reasons, you should set these settings back to the defaults after you've read this chapter.

The Search tab in the File Explorer Options applet controls the Search and Index feature in Windows 11. You can turn off using the search feature and change the behavior of non-indexed locations.

System Applet


The System applet in Control Panel is one of the most important applets, and technically it's not an applet. In most recent versions of Windows 11, the System applet in Control Panel will open the Settings app to the About screen. From within the Settings app panel, you can make a large number of configuration changes to a Windows machine. If you click Advanced System Settings, the classic System Properties (`sysdm.cpl`) will open. (See Figure 2.31 for the Windows 11 classic System Properties applet.) You can perform a number of functions in this applet, which can include some of the following options:

- Network Identification
- Device Manager
- Hardware
- User Profiles
- Environment Variables
- Startup and Recovery
- Performance

- System Protection
- Remote
- Computer Name
- Advanced

In the following sections, we will look more closely at the functionality of the tabs.

FIGURE 2.31 Windows System Properties Control Panel applet



Computer Name

This tab is used to define whether the machine is in a workgroup or a *domain* environment. We talk more about networking in Volume 2 Chapter 3, but in general terms, here's the difference between a workgroup and a domain:

Workgroup Loosely associated computers, each of which is its own security authority, that share a common *workgroup* name. Often used for 10 or fewer computers/workstations that do not require a server computer.

Domain A group of computers that are tightly connected or associated and that share a common domain name. Has a single authority (called a *domain controller*) that manages security for all the computers. Often used for 10 or more computers/workstations that connect to one or more server computers.


Hardware

This tab includes a number of tools that enable you to change how the hardware on your machine is used. The most useful is the ability to open the Device Manager directly from this tab. The other setting on this tab is how Windows behaves when you plug in devices. By default, Windows will automatically download *drivers*, *apps*, and custom icons from the devices plugged in. When you purchase a hardware device, odds are that it's been in the box for a while. By the time it gets made, packaged, stored, delivered to the store, stored again at the retailer, and then purchased by you, it's entirely likely that the company that made the device has updated the driver—even possibly a few times if there have been a lot of reported problems.

Advanced

The Advanced tab has several subheadings, each of which can be configured separately, as shown in Figure 2.32. The following options are among those on this tab.

FIGURE 2.32 Windows System Properties Advanced tab



Performance

Although it is hidden in the backwaters of Windows' system configuration settings, the Performance option holds some important settings that you may need to configure on a system. To access it, on the Advanced tab, click Settings in the Performance area.

In the Performance window, you can specify the size of your virtual memory and how the system handles the allocation of processor time. In Windows, you also use Performance to configure visual effects for the GUI.

How resources are allocated to the processor is normally not something that you will need to modify. It is set by default to optimize the system for foreground applications, making the system most responsive to the user who is running programs. This is generally best, but it means that any applications (databases, network services, and so on) that are run by the system are given less time by the system.




If the Windows machine will be working primarily as a network server, you may want to change the Performance option to Background Services. You can find this in the Performance Options on the Advanced tab. Otherwise, leave it as is.

Environment Variables

There are two types of *environment variables* (as shown in Figure 2.33), and you can access either one by clicking the Environment Variables button in the System Properties window.

FIGURE 2.33 Windows environment variables



User Variables Specify settings for an individual user and do not affect others who log into the machine.

System Variables Set for all users on the machine. System variables are used to provide information needed by the system when running applications or performing system tasks.




System and user variables were extremely important in the early days of DOS and Windows. Their importance has lessened with recent Windows versions, but this is the interface where TEMP variables, the location of the OS, and other important settings for Windows can be found.

User Profiles

In Windows, every user is automatically given a *user profile* when they log into the workstation. This profile contains information about their settings and preferences. Although it does not happen often, occasionally a user profile becomes corrupted or needs to be destroyed. Alternatively, if a particular profile is set up appropriately, you can copy it so that it is available for other users. To do either of these tasks, use the User Profiles settings to select the user profile with which you wish to work. You will be given three options, as shown in Figure 2.34.

FIGURE 2.34 Windows User Profiles Settings



Change Type The Change Type option enables you to configure a profile as local (the default) or roaming. If a user works at two machines, each machine will use a different profile. Updates to one machine will not be reflected on the other. If you have a network,

roaming profiles can be configured to allow a user to have a single profile anywhere on the network. Further discussion of this topic is beyond the scope of this book.


Delete The Delete option removes the user's profile entirely. When that user logs in again, they will be given a fresh profile taken from the system default. Any settings that they have added will be lost, as will any profile-related problems that they have caused.

Copy To The Copy To option copies a profile from one user to another. Often the source profile is a template set up to provide a standard configuration.

Startup And Recovery

The Windows Startup And Recovery options, shown in Figure 2.35, are relatively straightforward. They involve two areas: what to do during system startup and what to do in case of unexpected system shutdown.

FIGURE 2.35 Windows Startup And Recovery options



System Startup The System Startup option defaults to the Windows OS you installed, but you can change this default behavior if you like. Unless you are *dual-booting*, only one option is available; however, if you have another OS installed, you can change the Windows boot manager to load that as the default. You can also reduce the time the menu (time to display list of operating systems or recovery options) is displayed or remove the menu entirely.




If you choose to disable the menu completely on a dual-boot system, you will find that doing so may cause you annoyance in the future when you want to boot into a different OS but no longer have a choice to do so. Thus, you should always let the boot menu appear for at least 2–5 seconds if you are dual-booting.

System Failure A number of options are available in the Startup And Recovery dialog box for use in case of problems. These include writing an event about the problem, sending an alert to the network, and saving information about the problem to disk. These options come into play only in case of a major system problem.

System Protection

The System Protection tab lets you disable/enable and configure the System Restore feature, as shown in Figure 2.36. When System Restore is enabled on one or more drives, the operating system monitors the changes that you make on your drives. From time to time, it creates what is called a *restore point*. Then, if you have a system crash, it can restore your data back to the restore point. You can turn on System Restore for all drives on your system or for individual drives. Note that turning off System Restore on the system drive (the drive on which the OS is installed) automatically turns it off on all drives.


FIGURE 2.36 Windows System Protection Options



Remote

The Remote tab lets you enable or disable Remote Assistance and Remote Desktop, as shown in Figure 2.37. Remote Assistance permits people to access the system in response to requests issued by the local user using the Windows Remote Assistance tool. Remote Desktop permits people to log into the system at any time using the Remote Desktop Connection tool. This can help an administrator or other support person troubleshoot problems with the machine from a remote location.

FIGURE 2.37 Windows Remote options




Remote Assistance is enabled by default. It is handled at two levels. Having just Remote Assistance turned on allows the person connecting to view the computer's screen. To let that person take over the computer and be able to control the keyboard and mouse, click Advanced, and then, in the Remote Control section, click Allow This Computer To Be Controlled Remotely. You can also configure Remote Desktop here.

User Accounts Applet

The User Accounts applet allows you to view and create accounts for the Windows operating system. You can change the account name that appears on the Welcome and Start screen. You can also change the account type by selecting the Standard or Administrator radio button. In Windows 11, the Settings app for Accounts allows you to change the picture that is displayed along with your username.

In addition to the management of user accounts, the User Accounts applet allows you to change the *User Account Control (UAC)* settings for the operating system, as shown in Figure 2.38.


FIGURE 2.38 Windows User Accounts applet



Power Options Applet


The Power Options applet (`powercfg.cpl`) allows you to choose a power plan of Balanced and Power Saver, as shown in Figure 2.39. In addition to the Balanced and Power Saver plans, a system might contain CPU-specific power plans, High Performance, or Ultimate Performance, depending on the type of computer. Each power plan dictates when devices—namely, the display device and the computer—will turn off or be put to sleep.

FIGURE 2.39 Windows Power Options applet




When you click Change Plan Settings, you can change how fast the display is turned off and how fast the computer is put to sleep, as shown in Figure 2.40.

FIGURE 2.40 Windows Edit Plan Settings



Clicking Change Advanced Power Settings allows you to configure a number of settings based on power, as shown in Figure 2.41. These settings include specifying when the hard drive turns off, turning off the wireless adapter, specifying Internet options for JavaScript Timer Frequency, and determining the system cooling policy. The applet allows you to tweak your power policy, and you can always restore the plan defaults.

FIGURE 2.41 Windows Advanced Power Options settings



Power Management

The power plan configured in Windows will interface with the Advanced Configuration and Power Interface (ACPI). The ACPI must be supported by the system BIOS/UEFI in order to work properly. However, most computer hardware made in the last decade will support the ACPI. The ACPI on the computer hardware provides the operating system with the necessary methods for controlling the hardware. This is in contrast to Advanced Power Management (APM), which gave only a limited amount of power to the operating system and let the BIOS do all the real work. Because of this, it is not uncommon to find legacy systems that can support APM but not ACPI.

There are four main states of power management common in most operating systems:

Hibernate This state saves all the contents of memory to the hard drive, preserves all data and applications exactly where they are, and allows the computer to power off completely. When the system comes out of hibernation, it returns to its previous state by restoring the contents of memory from the hard drive. This power state consumes the least amount of power and the system will appear to be completely off in this state. This power state is usually only available on laptops and not desktops.

Sleep/Suspend In some operating systems, *Sleep* is used interchangeably with *Hibernate*, but that should not be the case. In Windows, Sleep puts the system in a low-power state, whereas Hibernate turns the system off—or at least it appears to be off.

Standby The *standby* state might appear to be identical to the sleep mode, as it restores the system quickly. However, it is not really a power state. Standby mode typically reduces power and turns the screen off.

Shutdown/Turn Off As the name implies, this shuts down everything and then turns off the system. You can shut down the system and have it stay off, or you can shut down the system and have it power back on (Restart). The latter option is often used when updates need to be applied.

Advanced Power Savings


If you are interested in saving power with a system that is not accessed often, one option is to employ Wake on LAN (WoL). Wake on LAN is an Ethernet standard implemented via a network interface card that allows a “sleeping” machine to awaken when it receives a wakeup signal. Wake on LAN cards have more problems than standard network cards. In our opinion, this is because they’re always on. In some cases, you’ll be unable to get the card working again unless you unplug the PC’s power supply and reset the card.

Choose What Closing The Lid Does

Another power-saving feature is to choose what happens when you close the lid of your laptop. This option is only available on laptop devices that have a lid closure sensor. The option can be found on the left-hand side of the Power Options Control Panel applet, and it’s labeled Choose What Closing The Lid Does. When the dialog box opens, you will see

drop-down menus for pressing the power button or sleep button and closing the lid, as shown in Figure 2.42. Depending on whether the laptop is plugged in or on battery power, you can choose a different option.

FIGURE 2.42 Windows Define Power Buttons And Turn On Password Protection dialog box



Fast Startup

Windows Fast Startup is another advanced feature that was originally introduced with Windows 8 as Fast Boot. The feature allows the system to hibernate during shutdown so that the system will appear to start up more quickly. The Fast Startup feature will attempt to know when a cold boot is required, such as when installing a program. If a cold boot is required, you have two options. The first option is to turn Fast Startup off in the Power Options System Settings, as shown in Figure 2.42. You can also use the shutdown `/s /t 0` command; the `/s` switch will shut down the system and the `/t 0` switch will do it immediately. The shutdown command also allows for shutdown of remote computers with the command `shutdown /m \\computername /s /t 0`.


USB Selective Suspend

The Universal Serial Bus (USB) selective suspend feature will allow the USB hub on the motherboard to suspend power to a device via the USB port. This is a handy feature to save battery power on a laptop if an external hard drive, a mouse, or some other device that requires power is not being used. However, this feature can also be problematic if you have a communication device connected via USB. You can turn this feature off on the Power Management tab of the device inside Device Manager, as shown in Figure 2.43.



There is a lot more going on than just powering off a USB port with the USB selective suspend feature. The operating system communicates to the device via the device driver and signals the device itself to power down. The system or the device can then send a wake-up signal to restore operations. You can find more information at <https://learn.microsoft.com/en-us/windows-hardware/drivers/usbcon/usb-selective-suspend>.


FIGURE 2.43 Device Manager, Device Properties detail, Power Management tab



Credential Manager Applet

The Credential Manager applet allows you to manage stored credentials for applications such as the web-based credentials and Microsoft Edge browsers, as well as the operating system itself. The Credential Manager service built into Microsoft operating systems stores username and password credentials in an encrypted database. The applet allows you to interact with the service to make changes to the stored credentials, as shown in Figure 2.44.


FIGURE 2.44 Windows Credential Manager applet



Programs and Features Applet

The Programs and Features applet (`appwiz.cpl`) allows you to view and uninstall desktop applications that are installed in Windows, as shown in Figure 2.45. You can also see the installed updates by clicking View Installed Updates on the left side of the applet.

The Programs and Features applet also allows you to install and remove features in Windows. When you click Turn Windows Features On Or Off on the left side, a dialog box will appear with a list of OS features that can be selected or deselected. Features are different from desktop applications because they are an integral part of the operating system. Examples include Hyper-V, Windows Subsystem for Linux, and Windows PowerShell 2.0, just to name a few.

FIGURE 2.45 Windows Programs and Features applet

Devices and Printers Applet

In Windows 11 the Devices and Printers Control Panel applet will now open the Bluetooth & Devices Settings screen. This Settings app will allow you to manage external devices such as external hard drives, printers, webcams, and any other external devices. From the Settings screen, you can discover and add new devices that are connected to your computer or network. Bluetooth & Devices is shown in Figure 2.46.

Sound Applet

The Sound applet (`mmsys.cpl`) allows you to view and change the default playback and recording device for sound on the system. By right-clicking the device and choosing Properties, you can view and modify the properties of a playback or recording device. The changes will vary by the device and the vendor of the device—common options are playback and recording levels and enhancements to the playback and recording levels. In addition, the Sound applet enables you to change the operating system's sound scheme, allowing you to change the various sounds, as shown in Figure 2.47.

Troubleshooting Applet

In Windows 11 the Troubleshooting applet will open into the Troubleshooting Settings screen. The Troubleshooter Settings screen does exactly what its name says it does—it allows troubleshooting of Programs, Hardware and Sound, Network and Internet, and System and Security. The Troubleshooter screen can be seen in Figure 2.48. Troubleshooting has one notable feature: application compatibility troubleshooting. When this feature identifies any application that is problematic, the Troubleshooting wizard will appear and Windows will try to troubleshoot the application compatibility.

FIGURE 2.46 Windows Bluetooth & Devices settings




FIGURE 2.47 Windows Sound applet





FIGURE 2.48 Windows Troubleshooting Settings app

Network and Sharing Center Applet

The Network and Sharing Center applet allows you to view and change the active network connections for the operating system, as shown in Figure 2.49. The applet displays the current network profile that the computer has been placed into by the Windows Firewall service. From the main page of the applet, you can click Change Adapter Settings to see the classic network adapter view, which allows the adapters to be configured manually. From the main page of the applet, you can also click Change Advanced Sharing Settings on the left side. This allows you to change network discovery options for the network profile chosen as well as turn on or off file and printer sharing globally for the operating system.

Device Manager Applet

As previously mentioned, the Device Manager applet (`hdwwiz.cpl`) was first introduced in Windows 95 and has hardly changed since its introduction over 20 years ago. If the applet looks familiar, it is the same as the MMC Device Manager (`devmgmt.msc`). Device Manager

allows you to view and change hardware devices on the operating system, as shown in Figure 2.50. This applet allows the administrator to load and update third-party drivers as well as drivers in the Windows Catalog. Device Manager is usually the first place you should check if a hardware device does not function after installation.

FIGURE 2.49 Windows Network and Sharing Center applet






FIGURE 2.50 Windows Device Manager



BitLocker Applet

BitLocker drive encryption is a low-level full-disk encryption feature that can be controlled from the BitLocker Drive Encryption applet, as shown in Figure 2.51. The BitLocker Drive Encryption applet also allows you to turn on BitLocker drive encryption on removable media devices.

FIGURE 2.51 Windows BitLocker Drive Encryption applet



Windows Defender Firewall

The Windows Firewall (`firewall.cpl`), which in Windows 11 is named Windows Defender Firewall, is used to block access from the network (be it internal or the Internet). While *host-based firewalls* are not as secure as other types of firewalls, this was a great move in the right direction. It first appeared in Windows XP Service Pack 2 but was released as a polished product in Windows Vista.

Figure 2.52 shows the opening screen of Windows Defender Firewall in Windows 11. Windows Defender Firewall is turned on by default. It also blocks incoming traffic by default.

Clicking Advanced Settings on the left opens the Windows Defender Firewall with Advanced Security MMC, as shown in Figure 2.53. This MMC allows great control of the Windows Defender Firewall controls for inbound and outbound rules. Inbound rules can be created to allow and deny network traffic inbound to applications and the operating system. Outbound rules can be created to allow and deny network traffic leaving the operating system. Only inbound firewall rules are restricted by default. When an application attempts to listen to specific network traffic, the operating system will display a dialog box asking you to confirm the network activity. This dialog box will automatically create an inbound rule

for the application's traffic. The Windows Defender Firewall with Advanced Security MMC allows you to pre-create the rule. In addition, you can create connection security rules for authenticating and encrypting traffic.

FIGURE 2.52 Windows Defender Firewall in Windows 11






FIGURE 2.53 Windows Defender Firewall with Advanced Security in Windows



Mail Applet

When Microsoft Outlook is installed in Windows, it is configured through the Mail applet. If you do not have Microsoft Outlook installed, the icon will simply not show up in Control Panel. The Mail applet is how you configure an email account, additional data files (OST files), RSS feeds, SharePoint lists, Internet calendars, published calendars, and address books, as shown in Figure 2.54. In the main dialog box you can even set up different profiles for different accounts. If more than one profile is configured, you will be prompted when Outlook launches to specify the profile you want to use.

FIGURE 2.54 Outlook Mail applet




Outlook will usually configure itself automatically by asking a series of questions when it first launches. The initial configuration will be stored and can be accessed by launching the Mail applet.

Indexing Options Applet


The Indexing service was introduced as a desktop search engine with Windows NT 4.0. Today the Indexing service is an integral part of Windows 11. It's an exceptional feature that is a requirement for today's volume of data. The Indexing service will systematically index files such as Microsoft Office documents, PDFs, text files, and many other files types. When searching for a word using File Explorer in a folder that is indexed, the Indexing service is queried directly, thus returning fast results. If the folder is not indexed, then the search process grinds through each file in the folder and produces results at a much slower pace.

You access the Indexing Options dialog box, shown in Figure 2.55, by using the Index Options applet in Control Panel. The default locations indexed are Start Menu and the Users folder (excluding AppData files). You can add locations to be indexed if you store files outside the normal Documents or Desktop locations that are contained inside the Users folder.

FIGURE 2.55 Indexing Options dialog box



By clicking Advanced in the Indexing Options dialog box, you open the dialog box shown in Figure 2.56. There you can choose to index encrypted files and to treat words with diacritics (accents) as different words. You can also rebuild the index in an attempt to fix missing documents from your search. This dialog box also allows you to relocate the index database. The File Types tab allows you to add various file types to index. It contains a very inclusive range of file types, but by default many are set to index just metadata on the file. Important file types like DOCX and PDF are set to index the contents.


FIGURE 2.56 Advanced Options for indexing

Ease of Access Center Applet


The Ease of Access Center applet contains various settings that make it easier to use Windows for motor and sensory impaired users, as shown in Figure 2.57. There is a wide range of tools, such as a magnifier, narrator, onscreen keyboard, and a high-contrast color scheme. In the many versions of Windows over the years, the accessibility tools have become expansive. Microsoft has continually added tools and settings to this applet to allow anyone with an impairment to use the operating system to its fullest.

Administrative Tools and Windows Tools Applet

In Windows 11 the Administrative Tools applet has been replaced by the Windows Tools applet in Control Panel. Actually, it isn't really an applet at all; it is like a shortcut to various tools, as shown in Figure 2.58. These tools all have a common theme: administering the operating system. Many of these tools can be accessed in other ways, such as right-clicking the Start button or using the Start menu and expanding Windows Administrative Tools in versions prior to Windows 11 or Windows Tools in Windows 11.

FIGURE 2.57 Ease of Access Center applet

In this array of tools, you can configure component services, perform disk cleanup and defragment the disk, set up iSCSI connections, edit the local security policy, set up Open Database Connectivity (ODBC) connectors, schedule tasks, and perform memory diagnostics. These are just a few tasks that we haven't already covered.


FIGURE 2.58 Windows 11 WindowsTools applet

Windows Settings

The Windows Settings app first made its debut in Windows 8. It was Microsoft's attempt to make configuring Windows simpler for end users. Many of the configuration tasks formerly performed in Control Panel have been either duplicated in the Settings apps or replaced entirely. The appearance of the Settings app has created anxiety for both end users and administrators alike, because it's a change (albeit unwanted) from the Control Panel that has been around since Windows 95.

You can open the Settings app by clicking the Start menu and selecting the gear icon with the title Settings. The Settings app will open to the screen shown in Figure 2.59. Here, you can search for the setting you need, or you can choose from various categories. The search capability has been a welcomed feature, since every release of Windows introduces new settings.

Many of the settings that are covered in this section have also been covered in the Control Panel section, as per the objectives of the CompTIA 220-1202 exam. It is good to know both ways to access settings, since many of the settings in Control Panel have not been moved over to the Settings app. Likewise, many of the settings in the Settings app can only be found in the app, because they are entirely new settings.

FIGURE 2.59 Windows 11 Settings app

Time & Language


The Time & Language screen in the Settings app allows you to change settings related to the date and time, language and region, keyboard preferences, and speech in Windows 11. Click Time & Language to change anything that you can change in Control Panel. The date and time can be changed, as shown in Figure 2.60.

The Set Time Automatically switch is set by default and synchronizes with Network Time Protocol (NTP) server of `time.windows.com`. You can choose to have Windows adjust the time zone automatically, or you can manually change the time zone. This dialog box also allows for customized calendars in the Taskbar.

In the upper-right corner, you can select Region or you can open Time & Language and select Language & Region. Changing the Country Or Region setting allows Windows to deliver content relevant to the area where you reside. This dialog box also allows you to


change the way values such as money, time, and date are formatted. If you select Windows Display Language, you can set the Windows display language as well as the preferred language, as shown in Figure 2.61. In addition, you can set the keyboard language and speech language preferences from the Time & Language menu.

FIGURE 2.60 Date & Time settings




Windows Update

Update and Security from Windows 10 has been replaced by Windows Update in Windows 11. The Windows Update screen is where you can access all settings related to Windows Updates and their delivery, as shown in Figure 2.62.

FIGURE 2.61 Language settings

From this initial screen, you can check for Windows Updates, as well as control downloads and installation, view the optional updates, pause updates, view update history, and set advanced options. Advanced options allow you to specify whether you want to receive updates for other Microsoft products, download updates over metered connections, change your active hours, determine how soon Windows is restarted after updates are applied, update notifications, and whether to pause updates until a certain date.

If you click Delivery Optimization on the Advanced Options screen for the Windows Update settings, you can change where Windows receives updates. By default, Windows will attempt to download updates from other PCs on the local network in order to conserve bandwidth. If the updates are not available, then Windows will download the updates from the Windows Update service.

FIGURE 2.62 Windows Update Settings

Windows Security

The Windows Security screen can be accessed via the Privacy & Security screen and selecting Windows Security. This will allow you to change a number of security-related items, as shown in Figure 2.63. From this screen, you can view Virus & Threat Protection, Account Protection, Firewall & Network Protection, App & Browser Control, Device Security, Device Performance & Health, and Family Options. The Windows Security screen is a one-stop shop for everything security-related for the operating system. You can click on each of the items and it will open in the respective settings for the item.

If you select Open Windows Security, as shown in Figure 2.63 a new standalone Windows Security screen will open that contains many more items that can be controlled. The standalone Windows Security settings screen can be seen in Figure 2.64.


FIGURE 2.63 Windows Security settings

Virus & Threat Protection The Virus & Threat Protection section provides you with the ability to initiate scans of the operating system, manage Virus & Threat protection settings, check for updates related to Virus & Threat protection, and set up Ransomware Protection, as shown in Figure 2.65.

Account Protection In this section you can sign into a Microsoft Account that allows for seamless synchronization of security settings and updates across multiple devices, ensuring consistent and up-to-date protection. Hello provides you with the ability to sign in with biometrics. Windows Dynamic Lock is a security feature that automatically locks your computer when a paired Bluetooth device, such as your smartphone,

moves out of range. This helps prevent unauthorized access if you step away from your PC and forget to lock it manually. These options can be found seen in Figure 2.66.

FIGURE 2.64 Stand-alone Windows Security settings




Firewall & Network Protection The Firewall & Network Protection screen in the Settings app allows you to view and manage your active firewall setting, as well as allowing apps through the firewall, as shown in Figure 2.67.


App & Browser Control The App & Browser Control screen protects you from untrusted applications and malicious websites. Smart App Control is available on clean installs of Windows 11; it starts in evaluation mode to see if you are a good candidate for the setting to be turned on. Once it is turned on, it will block programs that act

malicious or evasive. This setting is not recommended for developers or administrators, as some of the tweaks to the operating system might be blocked or perceived as malicious. Reputation-Based Protection will protect you from potentially unwanted application (PUA) by monitoring a program's reputation. In addition, you can control SmartScreen for Edge, which protects you from phishing sites and malicious downloads. Exploit Protection protects you from specific attacks and works in combination with Microsoft Defender subscriptions.

FIGURE 2.65 Virus & Threat Protection settings




Device Security The Device Security screen provides you with three options of Core Isolation, Security Processor, and Data Encryption. Core Isolation allows you to control memory integrity, which can be used to identify malicious software. Security Processor provides details on the Trusted Platform Module (TPM) specification. Data Encryption is a link that opens the BitLocker Control Panel applet.

FIGURE 2.66 Account Protection settings

Device Performance & Health This screen does not allow any controls of settings—it is simply a display of a health report for storage, apps and software, and the Windows time service, as shown in Figure 2.68. The screen also explains what Fresh Start is, and it provides a link to online content that offers details about the Device Performance & Health screen.

Family Options This screen does not allow any control of settings—it simply explains some of the parental controls you can employ. The screen also supplies a link to view family settings online. The online portal allows you to set up the controls for your family, such as content filters, screen time, spending, activity reporting, email, and calendar, just to name a few. The controls in the web page are outside the scope of the CompTIA exam, but we included this topic for completeness.

Protection History This screen does not allow any control of settings. It does provide the most recent security-related events, as shown in Figure 2.69.


FIGURE 2.67 Firewall & Network Protection settings

Personalization

The Personalization Settings screen allows you to personalize the entire operating system. The Personalization screen consolidated a number of other Control Panel applets when Windows 10 was first introduced. The new functionality introduced in Windows 10 was carried forward to Windows 11, shown in Figure 2.70.

Background Obviously, this section allows you to configure the background, formerly known as the wallpaper. You can choose a standard image, or choose to use your own photo. You can even choose Windows Spotlight for your background. The Windows Spotlight will download pictures from Bing. In addition, you can choose contrast themes for low vision or light sensitivity.


Colors The Colors section allows you to change the colors for the Windows controls and application controls. You can choose the transparency effects for Windows and the accent colors. An option to choose contrast themes for low vision or light sensitivity also exists.

FIGURE 2.68 Device Performance & Health

Themes As with prior versions of Windows, the theme can be changed using the Themes section. Changing the theme will change the background, colors, sounds, and mouse cursor. You can even download more themes from the Microsoft Store. The option to choose contrast themes for low vision or light sensitivity also exists.

Dynamic Lighting Dynamic Lighting allows you to change the RGB colors of connected devices so that third-party software is not needed. This section allows you to control a myriad of devices that have RGB LEDs.

Lock Screen Lock Screen allows you to configure how the lock screen looks and what is displayed on the screen when it is locked. By default, the lock screen is set to Windows Spotlight. The Spotlight feature downloads pictures from Bing and displays them on the lock screen. You can also configure which applications will display their status on the lock screen. This section also allows you to control the screen timeout settings related to power and the screen saver timeout.


FIGURE 2.69 Protection History

Text Input This section allows you to customize the text inputs prompts, the touch keyboard look, the voice typing prompt, and emojis. You can apply preinstalled themes or create a customized theme.

Start The Start section allows you to personalize the Start menu. You can change a number of settings, such as displaying the app list in the Start menu, showing recently added apps, and showing the suggestions, just to name a few settings.

Taskbar The Taskbar section allows you to change a number of settings. You can turn on or off Taskbar items, such as Search, Task View, and widgets. In this section, you can also control system tray icons for the pen menu and touch keyboard. By changing Taskbar Behaviors, you can automatically hide the Taskbar, and change the justification of the Taskbar from center to left. There are several other settings, such as combining Taskbar buttons into a single icon, flashing icons, and several other settings.


Fonts The Fonts section allows you to view all the installed fonts on Windows, as well as install new fonts by dragging and dropping them. The Fonts section also contains a link to open the Microsoft Store so you can download additional fonts.

FIGURE 2.70 Windows 11 Personalization settings

Apps

The Apps Settings screen controls every aspect about application on the operating system. This screen will eventually replace the Program and Features Control Panel applet, since it performs many of the same functions.

Installed Apps The Installed Apps section displays all the applications installed on the operating system, as shown in Figure 2.71. You can also uninstall apps by clicking the three dots on the right-hand side of the application section and selecting Uninstall. By clicking Advanced Options, you can choose to reset the application, repair the application, uninstall the application, and terminate the application. When you click the three dots on some applications, you can only choose to modify the application.

FIGURE 2.71 Installed Apps settings

Advanced App Settings The Advanced App Settings will allow you to change the source of your applications. The default is Anywhere, so you can continue to install downloaded applications. You can also choose the Microsoft Store Only option, or Microsoft Store Preferred if there is a comparable application. The Advanced App Settings also allow you to choose your sharing level for applications across your other devices. In addition, you can uninstall updates through the Advanced App Settings.

Default Apps In the Default Apps section, you can select the default app for email, maps, music, photos, videos, and websites. You also have the option to set the default app by the file type of the file or the network protocol being used. This is also known as the link type, which can be changed by selecting Choose Defaults By Link Type.

Offline Maps Offline Maps allows you to download maps, as well as update the already downloaded maps. You can select the storage location for offline maps, as well as control map downloads over a metered connection.

Apps for Websites Use the Apps For Websites section to associate an app with a website. Doing so causes the app to “spring into action” when a website is visited, making it a seamless experience.

Video Playback The Video Playback section allows you to tweak how a video will look when played back on Windows. You can choose to play video at a lower resolution to save network bandwidth. Other settings include high dynamic range (HDR) settings and a shortcut to System > Display And System > Power.

Startup Use the Startup section to specify which applications start in the background. By adjusting unwanted application starting when you log in, you can speed up your logon experience.

Privacy

The Windows Privacy Settings are located on the Privacy & Security Settings screen. The Privacy Settings are alongside of the Windows Security settings we previously covered. The Privacy Settings allow you to control all your privacy concerns with the Windows operating system, as shown in Figure 2.72.

Find My Device This section allows you to sign into a Microsoft account so that you can track the device if it is lost or stolen. This section is mainly focused at mobile devices.


General Here you can control your advertising ID, how websites make content decisions for you, the tracking of app launches, and suggested content settings.

Speech The Speech section allows you to control how your voice is used for speech recognition with Microsoft’s online speech recognition technology.

Inking & Typing Personalization The Inking & Typing Personalization section allows you to control whether your handwriting is used to build a personal dictionary of words.

Diagnostic & Feedback In Diagnostic & Feedback, you control how your personal information is used to provide diagnostic and feedback statistics to Microsoft. You can improve your inking and typing, and allow serving of tailored tips and ads. You can also view diagnostic data and delete the personalized data, such as the custom dictionary.

Activity History The Activity History section lets you control whether your activity is stored and tracked on the device. The history is not limited to recent documents—Windows can also track previous website visits. This allows you to continue where you left off.

FIGURE 2.72 Privacy settings

In prior versions of Windows, when you pressed the Windows key+Tab, Windows would show you all your previous documents that you had opened. This feature was called Windows Timeline. This feature was retired from later versions of Windows 11. Pressing the Windows key+Tab will still show you open windows.

Search Permissions The Search Permissions section allows you to control content moderations. You can filter out explicit content from web searches using the Strict setting, or turn it off completely; by default, it is set to Moderate. There are several other controls that allow you to search your personal, work, or school account for documents.


Searching Windows The Searching Windows section allows you to control how documents are searched on your computer. You can exclude folders from being included in searches. If you click the Advanced Indexing Options section, the Indexing Options Control Panel applet will open.

App Permissions In addition to the aforementioned Windows permission settings, you can view and control App permissions. This is the most powerful feature of the Privacy section, and it mirrors the controls found on mobile operating systems today. The various permissions that can be viewed and controlled are as follows: Location, Camera, Microphone, Voice Activation, Notifications, Account Info, Contacts, Calendar, Phone Calls, Call History, Email, Tasks, Messaging, Radios, Other Devices, Background Apps, App Diagnostics, Automatic File Downloads, Documents, Pictures, Videos, and File System. This is quite an exhaustive list, and you can review and control each of these permissions for a particular application.

System

The System section allows you to change a multitude of settings that pertain to the operating system, as shown in Figure 2.73. Just about every aspect of the operating system can be controlled.

FIGURE 2.73 System settings



Display In the Display section, you can arrange your monitors if you have more than one. You can also change how the additional monitors operate, such as extending or duplicating your desktop. You can also turn on the feature called Night Light, which restricts the blue light the display normally emits. The Display section also allows you to tune the Windows high dynamic range (HDR) of colors on your display. The most important settings are probably the display resolution and Scale And Layout settings, which allow you to get the most out of your display.

Sound In the Sound section, you select your output and input devices. You can also click the Troubleshoot button to help you identify sound issues. The Sound section is similar to the Sounds Control Panel applet, because it allows you to change the sound devices and control volume levels.

Notifications The Notifications section allows you to change the way the operating system notifications behave. You can control all operating system notifications, change lock screen notifications, control reminder and incoming VoIP calls on the lock screen, and specify whether notifications play sounds, among other settings.

Focus The Focus section allows you to control which notifications come to your attention and when they notify you; you can, for example, choose to suppress notifications when you are playing a video game.

Power The Power section is identical to the Power Control Panel applet. Here you can change when the screen turns off and when the operating system enters sleep mode. You can even change the power mode, similar to the Power Control Panel applet. You can also review recommendations to conserve energy.

Storage The Storage section provides a graphical overview of space used on the local disk. Clicking Installs Apps displays all the applications you can uninstall on the operating system. Temporary Files displays all the various temporary files on the operating system (you can then choose to remove them). This option is similar to the disk cleanup feature. Another feature called Storage Sense can be turned on, which automatically frees up space on the local disk by removing unneeded files.

Nearby Sharing The Nearby Sharing section allows you to control how your device is discovered by other devices. By default, the Nearby Sharing feature is turned off. You can turn the feature on for only your devices or for everyone. In addition, you can choose where to save received files, and you can change the discoverable name of your device.

Multitasking Use the Multitasking section to control Snap Assist, which is how an application or window snaps into a corner of the screen. You can also change the way the Alt + Tab keys display applications. In addition, you can configure how virtual desktops are used in Windows 11.

For Developers There are several controls in the For Developers section that only a developer can truly appreciate and only a developer should configure. Several of the settings in this section can make your PC insecure, as they share data on the local network

and to other applications. There are some notable features that even a network admin could appreciate, such as the File Explorer options. You can adjust whether hidden files or file extensions are shown. Set these settings at your own risk. The rule of thumb is unless you need it, don't set it.

Activation The Activation section allows you to see the current state of the operating system's activation. You can also change the product key in this section.

Troubleshooting This section allows you to control the behavior of the troubleshooter. By default, it is set to Don't Run Any. This setting can be changed to automatically run or ask you before running. In addition, if you click the Other Troubleshooters menu item, you will open another screen where you can run a multitude of troubleshooters.

Recovery The Recovery setting screen allows you to fix problems without resetting your PC by directing you to the Troubleshooters section. You can fix problems with your PC using Windows Update and by reinstalling your operating system. The most popular options are Reset This PC and Advanced Startup; both of these are covered Volume 2 Chapter 5, "Security Concepts."

Projecting To This PC The Projecting To This PC section allows you to control how other devices project their displays to Windows 11. The protocol used is called Miracast, which is a technology that allows screen sharing across devices.

Remote Desktop Use the Remote Desktop section to enable and disable the Remote Desktop feature, which allows you to connect remotely to a PC. We will cover these options in Volume 2 Chapter 8, "Scripting, Remote Access, and Artificial Intelligence."

Clipboard The Clipboard section allows you to control how the Clipboard operates. You can turn on features like Clipboard History, which enables you to have multiple items in your Clipboard. You can even sync Clipboards across multiple devices.

System Components The System Components section allows you to control the various components of the operating system, such as the Game Bar, Get Help, and Microsoft Store, just to name a few. You have similar options in resetting, repairing, and terminating the process as you do with installed apps.


Optional Features This section allows you to remove features, such as VBScript, Notepad, Print Management, and several others that are built in. In addition to removing, you can also install features like Hyper-V, .NET Framework, legacy components, and many others.

About The last section, About, allows you to view information about the PC and rename it, if you want. In addition, you can join an Active Directory domain or change workgroups. We will cover domains and workgroups in Volume 2 Chapter 3. There are several other options that will take you to other settings screens or Control Panel applets. These options include System Protection, Advanced System Settings, Product Key and Activation, Remote Desktop, Device Manager, and BitLocker.

Bluetooth & Devices


The Bluetooth & Devices Settings screen allows you to view, control, and configure all devices connected to the PC. If you remember back in the Control Panel section of this chapter, this is the settings screen that the applet is directed back to when you click the Devices And Printers Control Panel applet. The settings screen is shown in Figure 2.74.

FIGURE 2.74 Bluetooth & Devices settings



Devices The Devices section allows you to view and configure devices that are directly connected to the system, connected via Bluetooth, connected via another wireless technology, or wired directly, as shown in Figure 2.75. On this screen, you can add a Bluetooth, wireless display or dock, or any other device with a pairing process.

Printers & Scanners The Printers & Scanners section allows you to view and configure all of the installed printers and imaging devices on the operating system. You can allow Windows to manage the default printer, and Windows will select the most recently used printer as your default. You also have the option to add additional printers or scanners.

FIGURE 2.75 Devices settings

Mobile Devices This section allows you to add mobile devices that can be managed by the operating system. You can then configure the Phone Link options so you can receive calls on your PC. This feature can be set up for Android or iPhone.

Cameras This section allows you to search for network cameras on your network so that they can be used by the operating system. In addition, this is where you will find your webcam settings.

Mouse The Mouse section lets you change how the mouse behaves on the operating system. You can change settings like which button on the mouse is your primary button, your cursor speed, how the scroll wheel advances, and other mouse-related settings.

Pen & Windows Ink The Pen & Windows Ink section contains settings related to handwriting, such as the font you want to use when converting handwriting.


AutoPlay The AutoPlay section allows you to control how AutoPlay works for media and devices connected to Windows. You can choose a default action for removable drives or memory cards inserted.

USB The USB section allows you to control notifications if a USB device is not working correctly or has issues.

Network & Internet

Use the Network & Internet section to view and control network settings for Windows as well as Internet settings. The opening screen is the status section, and it displays the current network status for the Network connection, as shown in Figure 2.76. On this screen you can click Properties and then change various properties of the network connection, such as the public or private firewall posture, metered connection setting, and IP addressing. A newer feature lets you view the data usage for each application, which makes it easy to find an application that uses a lot of bandwidth.

FIGURE 2.76 Network & Internet settings



Ethernet or Wi-Fi The Ethernet or Wi-Fi section will change depending on your connection to the network. This section works as described earlier, mainly because the status section link redirects to this menu item.

VPN Use the VPN (virtual private network) section to view and configure settings for VPN connections. You can add a VPN connection or change advanced options, such as allowing VPN over metered networks or allowing VPN connections while roaming if a cellular modem is being used.

Proxy The Proxy Settings screen is where you can configure a proxy for the currently logged-on user. Here you set the proxy that Microsoft Edge will use. Other applications that use the common Microsoft web controls will also use the proxy server. By default, Automatically Detect Settings is enabled, but you can elect to use a setup script instead. It is a common task to manually set a proxy server, which will be in the form of an IP address or *fully qualified domain name (FQDN)*, along with the port. You also specify a bypass list of addresses that is allowed to bypass the proxy server.

Dial-Up Although it's unlikely you have a dial-up connection, the Network & Internet settings screen includes a section for configuring dial-up connections.

Advanced Network Settings The Advanced Network Settings screen allows you to configure a number of advanced settings for the network adapters. Use this section to configure advanced sharing options, such as network discovery and file and printer sharing. You can also reset the network connection and open the Firewall and Internet Options from this section.

Gaming



The Gaming settings screen was originally introduced with Windows 8 to create a seamless interface between the Xbox platform and the PC. Today, it has evolved into a very rich feature, and here you control how gaming is performed on Windows 11. The Gaming settings screen is shown in Figure 2.77.

Game Bar In this section, you can control whether the game bar is active during a game and how it launches. You can change any of the shortcut keys related to the Xbox Game Bar.

The most important key sequence is Windows key+ G, which launches the Xbox Game Bar, as shown in Figure 2.78. However, if you have an Xbox controller, the Xbox key will launch the game bar.

Captures The Captures section allows you to configure where screenshots and recorded captures are saved. You can manage all aspects of the capture in this section, such as recording length, recording audio, audio quality, microphone and system volume levels, recorded frames per second, and overall video quality.


Game Mode Use the Game Mode section to control the game mode, which turns off Windows updates so they don't interrupt gameplay. You can also adjust the quality of gameplay to deliver the best frame rate, and manually change the Graphics settings for performance of either desktop apps or Microsoft Store apps.

FIGURE 2.77 Gaming settings**FIGURE 2.78** Windows 10 Xbox Game Bar

Accounts

The Accounts section allows you to view and configure all the settings for your user account, as well as other accounts on the operating system. The setting screen is shown in Figure 2.79.

FIGURE 2.79 Accounts settings



Your Info The default screen is the Your Info screen, and it will display all of the information about your account, such as name, email address, and account type. You also have the option of managing your Microsoft account online.

Email & Accounts The Email & Accounts section enables you to add an email account that is used for email, calendar, and contact information. You can also change accounts used by other apps, such as the Microsoft Store app, that require a login.

Sign-in Options Use Sign-in Options to change the way you log into Windows. The Windows Hello feature is configured on this screen. Windows Hello allows you to substitute your face, fingerprint, PIN, security key, or picture password for your actual password. The Hello feature works by storing your real credentials, such as your username and password, in Credential Manager. Credential Manager is then locked with this process. When you attempt to log in with a picture of your face, Windows Hello will unlock the credentials stored in Credential Manager and pass the actual username and password to the operating system. Dynamic Lock is another feature that can be configured in this section. Dynamic Lock will dynamically lock your computer when you walk away with a device that is paired to the laptop, such as a mobile device.

Family This screen does not allow any control of settings; it simply explains some of the parental controls you can employ. The screen also supplies a link to view family settings online. This is the same screen that can be found in the Windows Security settings screen under Family Options.

Other Users This screen allows you to set up other user accounts for the operating system. You can also set up the operating system to act as a kiosk. Uses might be a digital sign, interactive display, or something else where you don't need someone logged into the desktop. You launch a wizard that will create a local account, and you can then choose a kiosk app. When the kiosk mode is enabled, the operating system will boot up and automatically log in as the local user created, and the configured app will run.

Windows Backup The Backup section allows you to back up files to OneDrive. You can remember the apps on the current operating system. This allows you to restore both the apps and the data when you log into a new computer. This feature requires a Microsoft account and enough storage to back up your files.


Access Work Or School The Access Work Or School section is used to connect the operating system with a corporate or school account. These accounts usually contain mobile device management (MDM) settings. This passes some or all of the control of the operating system to the organization responsible for the account. The enrollment of the operating system into an MDM system can be performed with a provisioning package to help simplify the enrollment process. You can also export management log files for analysis if something is not functioning properly with the MDM control. You can set up an account for test taking, which locks down the operating system when it is logged into.

Passkeys The Passkeys section provides you with a way to manage the various passkeys for websites and applications. You can then use biometrics or a PIN through Windows Hello to unlock and give the application the key.

Accessibility

The Accessibility settings screen is almost identical to the Ease of Access Control Panel applet. In the future, we will probably see the Ease of Access Control Panel applet converged into the Accessibility Settings screen, like many other applets. The Accessibility Settings screen is divided into three categories—Vision, Hearing, and Interactions, as shown in Figure 2.80.

FIGURE 2.80 Accessibility settings



Text Size The Text Size screen allows you to adjust the size of the text displayed on the screen. This feature is particularly useful for individuals with visual impairments, as it enables them to increase the readability of text across the operating system and

applications. By modifying the text size, users can ensure that all textual content, including menus, dialog boxes, and system messages, is easily legible.

Visual Effects The Visual Effects setting provides options to customize the appearance and behavior of visual elements in Windows 11. You can enable or disable animations, transparency effects, and other visual enhancements to improve system performance or reduce visual clutter. This setting is beneficial for users who experience motion sensitivity or who prefer a more streamlined interface.

Mouse Pointer and Touch This setting allows you to customize the appearance and behavior of the mouse pointer and touch input. You can change the pointer size, color, and speed, as well as enable features like pointer trails and touch feedback. These adjustments help individuals with motor impairments or visual difficulties to interact more effectively with their devices.

Text Cursor The Text Cursor setting enables you to modify the appearance and behavior of the text cursor, also known as the caret. You can change the cursor's thickness, color, and blink rate to make it more visible and easier to locate within text fields. This feature is particularly useful for individuals with visual impairments or those who have difficulty tracking the cursor's movement.

Magnifier The Magnifier is a screen magnification tool that allows you to enlarge portions of the screen for better visibility. It offers various modes, including full-screen, lens, and docked views, to accommodate different user preferences. The Magnifier also provides options for adjusting the zoom level, following the mouse pointer, and inverting colors to enhance contrast.

Color Filter The Color Filter setting enables you to apply color filters to the screen to improve visibility and differentiate between colors. This feature is particularly helpful for individuals with color blindness or other visual impairments. Users can choose from various filter options, such as grayscale, inverted, and specific color filters, to suit their needs.

Contrast Themes Contrast Themes are high-contrast color schemes designed to make text and interface elements more distinguishable. These themes use a limited color palette with high contrast between foreground and background elements, making it easier for users with low vision or color blindness to navigate and interact with the system.

Narrator The Narrator is a screen reader built into Windows 11 that provides spoken feedback for onscreen content. It reads aloud text, buttons, and other interface elements, allowing users with visual impairments to navigate and interact with their devices. The Narrator offers various customization options, including voice selection, verbosity levels, and keyboard shortcuts.

Audio The Audio settings allows you to adjust sound-related accessibility features, such as mono audio, audio alerts, and sound schemes. These options help individuals with hearing impairments or auditory processing difficulties to better perceive and respond to audio cues within the operating system and applications.

Hearing Devices This setting provides support for connecting and configuring hearing devices, such as hearing aids and cochlear implants, to Windows 11. You can pair their devices via Bluetooth or other compatible methods and adjust settings to optimize audio output for their specific hearing needs.

Captions The Captions setting allows you to customize the appearance of closed captions and subtitles displayed in videos and other media. You can change the font, color, size, and background of captions to improve readability and ensure that the text is easily distinguishable from the video content.

Speech The Speech setting enables user to control their devices using voice commands and dictation. This feature is particularly useful for individuals with mobility impairments or those who prefer hands-free interaction. Users can configure speech recognition settings, train the system to recognize their voice, and customize voice command preferences.

Keyboard The Keyboard setting provides options to customize keyboard behavior and accessibility features. You can enable features like Sticky Keys, Filter Keys, and Toggle Keys to assist with typing and keyboard navigation. These settings are beneficial for individuals with motor impairments or those who have difficulty using a standard keyboard.

Mouse The Mouse setting allows you to adjust the behavior and appearance of the mouse pointer. You can change the pointer speed, enable pointer precision, and configure additional mouse buttons. These adjustments help individuals with motor impairments or those who require a more customized pointing experience.

Eye Control Eye Control is an accessibility feature that enables users to control their devices using eye-tracking technology. This setting allows individuals with severe mobility impairments to navigate the operating system, type text, and interact with applications using only their eye movements. Users can calibrate the eye-tracking device and customize settings to optimize the experience.

The Windows Registry

Windows configuration information is stored in a special configuration database known as the *Registry*. This centralized database contains environmental settings for various Windows programs. It also contains registration information that details which types of filename extensions are associated with which applications. So, when you double-click a file in Windows File Explorer, the associated application runs and opens the file that you double-clicked.



Exam Objective 1.4, “Given a scenario, use Microsoft Windows operating system features and tools,” requires you to know tools such as Registry Editor (`regedit.exe`).

The Registry was introduced with Windows 95. Most operating systems up until Windows 95 were configured through text files, which could be edited with almost any text editor. However, the Registry database is contained in a special binary file that can be edited only with the Registry Editor provided with Windows. The Registry Editor can be accessed a few different ways. You can launch it from the icon named Registry Editor in Control Panel > Windows Tools > Registry Editor. You can also launch it from the Start Menu > All > Windows Tools > Registry Editor. The fastest way is to open the Run box with Windows key+R and type **regedit.exe**.



Current versions of Windows have what appear to be two applications that can be used to edit the Registry: `regedit.exe` and `regedt32.exe` (with no i). In reality, however, `regedt32.exe` opens `regedit`. In Windows 11, they are the same application. In prior Windows versions they work similarly, but each had slightly different options for navigation and browsing.

The Registry is broken down into a series of separate areas called *hives*. The keys in each hive are divided into two basic sections—user settings and computer settings. In Windows, a number of files are created corresponding to each of the different hives. The names of most of these files do not have extensions, and their names are SYSTEM, SOFTWARE, SECURITY, SAM, and DEFAULT. One additional file, whose name does have an extension, is NTUSER.DAT.

The basic hives of the Registry are as follows:

HKEY_CLASSES_ROOT Includes information about which filename extensions map to particular applications.

HKEY_CURRENT_USER Holds all configuration information specific to a particular user, such as their desktop settings and history information.

HKEY_LOCAL_MACHINE Includes nearly all configuration information about the actual computer hardware and software.

HKEY_USERS Includes information about all users who have logged into the system. The HKEY_CURRENT_USER hive is actually a subkey of this hive.

HKEY_CURRENT_CONFIG Provides quick access to a number of commonly needed keys that are otherwise buried deep in the HKEY_LOCAL_MACHINE structure.

Modifying a Registry Entry

If you need to modify the Registry, you can modify the values in the database or create new entries or keys. You will find the options for adding a new element to the Registry on the

Edit menu. To edit an existing value, double-click the entry and modify it as needed. You need administrator-level access to modify the Registry.



Windows uses the Registry extensively to store all kinds of information. Indeed, the Registry holds most, if not all, of the configuration information for Windows. Modifying the Registry in Windows is a potentially dangerous task. Control Panel and other configuration tools are provided so that you have graphical tools for modifying system settings. Directly modifying the Registry can have unforeseen—and unpleasant—results. You should modify the Registry only when told to do so by an extremely trustworthy source or if you are absolutely certain that you have the knowledge to do so without wreaking havoc.

Restoring the Registry

Windows stores Registry information in several files on the hard drive. With Windows 10 later versions, you need to restore the files from a backup for the `systemroot\repair` directory by using the Windows Backup program. Repairing the Registry from a backup overwrote the Registry files in `systemroot\system32\config`.



The Last Known Good Configuration option no longer exists at boot time in Windows 10/11. If the operating system suffers two consecutive failed attempts to boot into Windows, the Windows Recovery Environment (WinRE) will automatically boot. From WinRE, the operating system will try to automatically fix common issues. However, the only way to recover a blue screen on boot is to choose Reset This PC in the WinRE and select Keep My Files. More information on WinRE can be found here:

<https://learn.microsoft.com/en-us/troubleshoot/windows-server/performance/use-winre-to-troubleshoot-startup-issue>



Real World Scenario

Beware of Editing the Registry

Just in case it hasn't sunk in yet: Be careful editing the Registry! There is no Undo button, nor do you have the safety net of choosing not to save your edits before you close. Once you make the change, it's made—for better or for worse.

There have been countless examples throughout our careers of people going in to edit the Registry without really knowing what they were doing. In many cases, making small

changes to the Registry, without having a viable backup, means having to reinstall Windows. At the very least, this is inconvenient.

Windows can help in this regard if you are in a networked environment with Windows-based servers. You can create system policies that prevent users from performing certain tasks, and the most important task to restrict is running Registry editors.

Remember that a system restore will restore the Registry to the state it was in when a restore point was saved. As a very last-resort option for system recovery, Windows uses the Windows Recovery Environment (WinRE) to do a complete PC reset. It is your goal to make sure that you never need to use this.

Disk Management

Where there are files, there are disks. That is, all the files and programs that we've talked about so far reside on disks. Disks are physical storage devices, and they also need to be managed. There are several aspects to disk management. One is concerned with getting disks ready to be able to store files and programs; another deals with backing up your data; and yet another involves checking the health of disks and optimizing their performance. We'll look at these aspects in more detail.



Exam Objective 1.1, "Explain common operating system (OS) types and their purposes," requires you to know the various filesystem types and disk partitioning. Be sure to know the differences between each of the filesystem types.

Getting Disks Ready to Store Files and Programs

In order for a hard disk to be able to hold files and programs, it has to be partitioned and formatted. *Partitioning* is the process of creating logical divisions on a hard drive. A hard drive can have one or more partitions. *Formatting* is the process of creating and configuring a file allocation table (FAT) and creating the root directory. The *New Technology Filesystem (NTFS)* is available with all the versions of Windows you need to know about for the exam, but others are also recognized and supported. The file table for NTFS is called the *Master File Table (MFT)*.

The following is a list of the major filesystems that are on the CompTIA objectives. They are split into three categories of Linux, Apple, and Windows.

Linux Filesystems

Linux filesystems are an essential component of the Linux operating system, providing a structured way to store, organize, and manage data on disk drives or partitions. They define how

data is stored, accessed, and retrieved, ensuring efficient data management and integrity. Linux supports a wide variety of filesystems, each with its own unique features and use cases. The following list is from the CompTIA objectives, but there are far more filesystems than these.

Extended File Allocation Table Created by Microsoft, *Extended File Allocation Table* (*exFAT*) is a proprietary filesystem of choice for flash drives where NTFS cannot be used (because of overhead) and FAT32 is not acceptable (due to filesystem limitations). Although exFAT is a proprietary filesystem, it is being adopted by third-party device vendors to overcome the filesystem limitations of FAT32. It is ideal for SD cards that hold a lot of information, and it is supported in all current versions of Windows. exFAT is the default filesystem on USB flash drives over 32 GB in size.

Network File System Network File System (NFS) was created by Sun Microsystems a number of decades ago and is widely used in UNIX and Linux environments. Starting with Windows 7, Microsoft includes support for NFS with Windows (but it is not always installed). Windows versions prior to Windows 7 (Vista) can install Windows Services for UNIX (SFU) to gain this support.

Extended File System The Extended File System (XFS) is a high-performance 64-bit journaling filesystem designed to handle large files and high-throughput environments. It divides the storage space into smaller sections, which helps in managing data more efficiently and allows for faster data access. XFS also has features that reduce data fragmentation, meaning it keeps files stored in a more organized way. It ensures data integrity by keeping track of changes, which helps in quick recovery if something goes wrong. Windows does not natively support the XFS filesystem.

Fourth Extended File System The name might sound the same as the Extended File System (XFS), but the ext filesystem is not derived from the XFS filesystem; each of these filesystems is unique. The ext3 (Third Extended File System) became the default filesystem for many distributions of Linux due to its journaling capabilities. Enhancements were added to it to create *ext4* (Fourth Extended File System), which is used with Android and other operating systems. It is possible to use ext4 with Windows, but doing so requires making some sacrifices in features for the sake of compatibility.

Swap Partition Swap partitions, which are found only within Linux and UNIX installations, are the equivalent of the page file (`pagefile.sys`) in Windows, except that they are their own partition type. They are used for *virtual memory* when the physical memory is exhausted.

Apple Filesystems

Apple filesystems are a fundamental part of the macOS, iOS, iPadOS, watchOS, and tvOS operating systems, providing a robust framework for organizing, storing, and managing data on various storage devices. The following are two of the main filesystems you may encounter on Apple devices.

Apple File System *Apple File System (APFS)* is a proprietary filesystem that was developed by Apple to replace HFS and HFS+. The filesystem was mainly developed to introduce new features with *solid-state storage drives (SSDs)*, such as compression, encryption, and snapshots, just to name a few. APFS also fixed some core issues with the HFS+ filesystem, such as the number of files and data integrity. The filesystem is the default filesystem for macOS Sierra 10.12.4, iOS 10.3, and tvOS 10.2 and later. Windows operating systems cannot read drives formatted in APFS.

Windows Filesystems

Windows filesystems are integral to the Windows operating system, providing a framework for organizing, storing, and managing data on various storage devices. The following are the most common files systems you may see as a technician.

File Allocation Table *FAT* is the acronym for the file table used to keep track of where files are within the filesystem. It's also the name given to this type of filesystem, introduced in 1981. The filesystems for many OSs have been built on the design of FAT but without its limitations. A FAT filesystem uses the *8.3 naming convention* (eight letters for the name, a period, and then a three-letter file identifier). This later became known as *FAT16* (to differentiate it from FAT32) because it used a 16-bit binary number to hold cluster-numbering information. Because of that number, the largest FAT disk partition that could be created was approximately 2 GB.

FAT32 The *FAT32* filesystem was introduced along with Windows 95 OEM Service Release 2. As disk sizes grew, so did the need to be able to format a partition larger than 2 GB. FAT32 was based more on VFAT than on FAT16. It allowed for 32-bit cluster addressing, which, in turn, provided for a maximum partition size of 2 terabytes (2,048 GB). It also included smaller cluster sizes to avoid wasted space. FAT32 support is included in current Windows versions.

New Technology Filesystem Introduced along with Windows NT, *New Technology Filesystem (NTFS)* is available with all current versions of Windows. NTFS is a much more advanced filesystem in almost every way than all versions of the FAT filesystem. It includes such features as individual file security encryption and compression, disk quotas, and RAID support, as well as support for extremely large file and partition sizes and disk transaction monitoring. It is the filesystem of choice for higher-performance computing.

Resilient File System The Resilient File System (ReFS) is a modern filesystem developed by Microsoft, designed to maximize data availability, scalability, and integrity. ReFS was developed by Microsoft to keep your data safe and reliable. It uses special checks to detect and fix any data corruption, ensuring your files stay intact. It can handle large amounts of data and supports long filenames, making it great for big projects. Additionally, ReFS regularly scans the storage to find and fix any hidden issues, keeping your system running smoothly.

Compact Disc File System Although not a filesystem that can be used on a hard drive, *Compact Disc File System* (CDFS) is the filesystem of choice for CD media. It has been used with 32-bit Windows versions since Windows 95. A CD mounted with the CDFS driver appears as a collection.

When you're installing Windows 11, the installer defaults to NTFS. However, you can use the partitioning tool to format the partition with exFAT. It's really unnecessary, since NTFS is a much better filesystem with respect to stability, avoiding corruption, and security. When you are formatting a storage drive, you can format it with FAT or NTFS. Storage devices formatted in FAT can be read by other operating systems, but when they are formatted in NTFS the operating system must explicitly support NTFS.

To format a partition from the command line, use the `format` command, which is available with all versions of Windows. You can run `format` from a command prompt or by right-clicking a drive in Windows File Explorer and selecting **Format**. However, when you install Windows, it performs the process of partitioning and formatting for you if a partitioned and formatted drive does not already exist.

You can usually choose between a *quick format* and a *full format*:

- With a quick format, a new file table is created on the hard disk, but files are not fully overwritten or erased from the disk.
- A quick format is much faster since the normal format fully erases any and all existing data on the hard disk (a time-consuming process).



Be extremely careful with the `format` command! When you format a drive, treasured data can be lost. Before using the `format` command, make sure that you back up the data you want to save.

Disk Management MMC

In Windows, you can manage hard drives through the Disk Management console. To access Disk Management in Windows 11, right-click the Start menu and then click **Disk Management**. Alternatively, you could start the Disk Management MMC by typing **diskmgmt.msc** in the Run dialog box (Windows key+R).


The Disk Management screen lets you view a lot of information about all the drives installed in your system, including CD/DVD-ROM drives (see Figure 2.81).

Given a scenario, there are many instances in which an administrator would turn to Disk Management during the course of trying to find the right storage solution. The Disk Management administrative console allows you to review all the logical partitions (volumes) configured and physical drives (disks) that are connected to the computer.

The logical volume is displayed in the upper portion of the Disk Management console. This view allows you to see at a glance all the configured volumes on the operating system.

Details about the volume letter or type, the layout type, the type of disk (basic vs. *dynamic*), filesystem, status, capacity, and free space can be seen in this logical view.

FIGURE 2.81 Disk Management MMC



The physical layout, displayed below the logical view, allows you to see the layout of the physical disks in a graphical format. Using both views, you can see the layout of a physical disk and its corresponding logical layout. The physical layout allows you to view the status of the drive, the size of the drive, and its health, in addition to the layout of volumes on the physical disk.

Windows supports three types of partition styles: Basic, GPT, and Dynamic. Partition styles are also known as partition schemes; the terms are used to describe the underlying structure of the partitioning of the physical disk.

Basic Disks *Master Boot Record (MBR)* is the standard BIOS partition table used and predates GPT. MBR disk storage can have a maximum of four primary partitions or three primary partitions and one extended partition containing multiple logical partitions, for a maximum of 26 logical partitions. The partition from which the operating

system boots must be designated as active and must be a primary partition. Only one partition on a disk can be marked active.

GPT Disks *GUID [globally unique identifier] Partition Table (GPT)* is the current partition table specification used with *Unified Extensible Firmware Interface (UEFI)* BIOS. GPT partition tables allocate 64-bit logical block addressing for larger partitions. If you need to create partitions 2 TB or larger, it is required. A GPT disk can contain up to 128 primary partitions. The large number of primary partitions supported eliminates the need to create extended/logical partitions.

Dynamic Disks *Dynamic disks* are proprietary to Windows operating systems. Dynamic disks can contain up to 2,000 volumes. All the dynamic disks installed in the operating system contain an identical 1 MB database (partition table) that describes all the other disks installed. The true benefit of this 1 MB database is moving fault-tolerant disks to another Windows operating system. The operating system will automatically recognize that the disks work together, and the *Logical Disk Management (LDM)* will configure itself according to the layout in the 1 MB database.



The terms *partition* and *volume* are interchangeable; they both mean the same thing. They are logical subdivisions of a physical disk. Dynamic disks use the terminology of volumes, and MBR and GPT use the terminology of partitions. Sometimes partitions will be referred to as volumes, such as when you use the diskpart command-line utility.

Right-clicking any volume opens a context menu that allows you to change the drive letter or paths, format, extend, shrink, delete, or add a mirror. Right-clicking any drive opens a context menu that allows you to create a new spanned, striped, mirrored, or RAID-5 volume. You can also convert to a dynamic or GPT disk.

Let's discuss the features and functions in further detail:

Initializing Disks The phrase *initializing a disk* means different things depending on if it is a hardware initialization or, in this case, a Windows initialization. When you initialize a disk in Windows, you are allowing the LDM to recognize the disk. When you initialize a disk in hardware disk systems, you are erasing the partition table and even the data. To initialize a disk, right-click the drive and select Initialize. This needs to be performed only once.

Mounting Partitions Once a partition (volume) is created, you need to make it available to the operating system for use. You have two basic options. The first option is to assign the volume a drive letter. The second option is to mount the volume in an empty NTFS folder. So you can mount an entire volume in a specific folder. You can perform this task by right-clicking a partition and selecting Change Drive Letter And Paths.

Extending Partitions When you install an operating system on a disk, you will more than likely use the entire disk for the installation of the operating system. However, if you ever fill it all up and buy a bigger drive, when you clone it over to the new larger disk, you will find that the partition is the same size. This is where extending a partition comes in handy. It is also very common for virtualized operating systems, since adding space to a virtual disk can be done on the fly. As long as there is unallocated space on the disk, you can right-click the partition (volume) and choose Extend Volume.

Shrinking Partitions Shrinking a partition, also known as splitting a partition, is not a common task, but it is done from time to time. It is usually performed when an installation is not using a fraction of the physical disk storage and another partition is required. You can perform this action by right-clicking the partition (volume) and choosing Shrink Volume.



Extending partitions and shrinking partitions require that you perform the task on the last partition (volume) created on the drive. This is not always the case. These functions will not always work from within the operating system. An open source tool called SystemRescueCD contains a tool called GParted. This tool allows you to reconfigure the disk layout and move partitions around, as well as shrink and extend partitions.

Adding Drives When you add a disk to the system, you normally have to power the computer down to attach the disk. However, servers have the capability to add drives on the fly, called *hot swap disks*. When a disk is added to the operating system, the operating system needs to know it. You can make the operating system take inventory of its disks by selecting Rescan Disks from the Action menu in the Disk Management console.

Adding Arrays Arrays of disks are multiple disks working together for a specific purpose. These are also known as a *Redundant Array of Independent (or Inexpensive) Disks (RAID)*. They are used for fault tolerance so that in the event of a disk failure, the system will continue to function. They can also be used to add performance to a system by using multiple disks. Striped disks, also known as *RAID-0*, require two or more disks and are used for performance; they do not offer any fault tolerance. A mirror, also known as *RAID-1*, is the most basic type of fault tolerance; it creates an identical mirror of the information on another disk. Striped disks with parity, also known as *RAID-5*, offer performance and fault tolerance and require three or more disks. Windows is unique in that it allows volumes to participate in array groups, rather than requiring the entire disk. A combination of mirroring (*RAID-1*) and striping (*RAID-0*) can be used together as *RAID-10*. *RAID-10* offers both fault tolerance of a mirrored disk and the performance of striping across the mirrors. *RAID-10* requires a minimum of four disks to operate.




Windows can create RAID groups of 0, 1, and 5, but RAID 10 is not directly supported. There are various tricks to making a RAID 10 volume, such as creating two mirrors and then striping across the mirrors. However, this adds a level of complexity, and we recommend that if you need RAID 10 you purchase a hardware solution such as a dedicated RAID card. Volume 1 Chapter 2, “Expansion Cards, Storage Devices, and Power Supplies,” provided more information about the hardware implementations of RAID.

Storage Spaces This Windows feature was initially introduced in Windows Server 2012 and Windows 8. It allows for a group of drives to be placed into a pool of storage that can be configured for fault tolerance. The benefit is that more disks can be added later to extend a storage pool dynamically. The one caveat is that the hard drives can either be managed by the Disk Management MMC or Storage Spaces, but not both.

The *Storage Spaces* feature lets you manage a variety of hard drives. A unique feature is its capability to create a storage pool using two or more external (USB-attached) hard drives. Storage Spaces is not exclusive to external hard drives; internal hard drives can be used as well.

In Storage Spaces, you can create a two-way mirror, a three-way mirror, or parity resiliency of your data. A two-way mirror is identical in functionality to a *RAID-1* mirror. A three-way mirror is similar to a RAID-1 mirror in that it duplicates the data on two other drives; unlike RAID-1, however, you can lose two hard drives and retain data. A parity resiliency type is identical in functionality to a *RAID-5* (striping with parity). You can also use Storage Spaces without resiliency, which is called a *simple volume*, as shown in Figure 2.82.

FIGURE 2.82 Storage Spaces



Checking the Health of Hard Disks and Optimizing Their Performance



Exam Objective 1.4, “Given a scenario, use Microsoft Windows operating system features and tools,” requires you to know tools such as Disk Defragment (dfrgui.exe).

As time goes on, it’s important to check the health of Windows computers’ hard disks and to optimize their performance. Windows provides you with several tools to do so, some of which we’ve already mentioned in this chapter. One important tool is Disk Defragmenter, which has existed in almost all versions of Windows.

When files are written to a hard drive, they’re not always written contiguously, or with all the data in a single location. Files are stored on the disk in numbered blocks, similar to PO boxes. When they are written, they are written to free blocks. As a result, file data is spread out over the disk, and the time it takes to retrieve files from the disk increases. Defragmenting a disk involves analyzing the disk and then consolidating fragmented files and folders so that they occupy a contiguous space (consecutive blocks). This increases performance during file retrieval, since the hard drive arm in mechanical hard disks needs to travel less to retrieve the blocks. Defragmentation of the filesystem is not required on solid-state drives (SSDs), since the fragmentation of data blocks does not slow down the retrieval of the memory locations.

To access Disk Defragmenter, follow these steps:

1. Click Start > All > Windows Tools > Defragment And Optimize Disk. Or you can click the Start menu and type **defragment**, then select it.
2. In the list of drives, select the drive that you want to defragment, and then click Analyze. When the analysis is finished, Disk Defragmenter tells you how much the drive is fragmented and whether defragmentation is recommended.
3. If defragmentation is recommended, click Defragment.

Be aware that for large disks with a lot of fragmented files, this process can take quite some time to finish.



In Windows, you can also access Disk Defragmenter through the properties of any partition listed in Disk Management (or, even easier, any disk root under This PC in Windows File Explorer). Select the Tools tab, and then click Defragment.

Summary

In this chapter, you learned about some of the tools that can be used with Windows. We covered basic Windows management concepts, including managing disks, using filesystems, and understanding directory structure. Keeping your computer healthy will save you a lot of stress.

With the basic knowledge gained in this chapter, you are now ready to learn how to interact with the most popular operating systems in use today. These topics are covered in the next four chapters.

Exam Essentials

Know the various filesystems and the differences between them. The most commonly used filesystem on Windows hard drives is NTFS. FAT32 is older and perhaps a bit quicker for smaller hard drives and USB thumb drives, but NTFS adds a plethora of important features, including security and auditing. exFAT is a filesystem Microsoft adopted to address large filesystems that FAT32 had problems addressing. The Extended filesystem (XFS) is a high-performance 64-bit journaling file system designed to handle large files and high-throughput environments and used with Linux. The Resilient File System (ReFS) is a Windows filesystem designed to maximize data availability, scalability, and integrity. Additionally, you should know that the ext3, ext4, and APFS filesystems are not compatible with Windows without third-party software.

Understand the various partitioning for Windows installs. The GUID Partition Table (GPT) is a modern specification used with UEFI BIOS, allowing for larger partitions and supporting up to 128 primary partitions. The Master Boot Record (MBR) is the traditional BIOS partition table, allowing for a maximum of four primary partitions or three primary partitions and one extended partition. Formatting a drive involves creating a file allocation table and configuring the root directory, with NTFS being the most commonly used filesystem.

Know which Windows Control Panel utility to use for a given scenario. A number of Control Panel applets allow you to change Windows' functionality. With the introduction of Windows 10, some of the functionality is split between the Settings app and the classic Control Panel applets; this is carried forward with Windows 11. You should be proficient with both for a given scenario in which changes to the operating system need to be made.

Know which Windows Settings section is used for a given scenario. The Settings app is slowly replacing Control Panel as the main method of configuring Windows. You should know the major categories of the Settings app and the various settings each of them contains.

Know the main administrative tools. You should know the primary graphical tools for troubleshooting Windows and configuring the operating system. These include the disk management tools, Windows Tools, Device Manager, Task Manager, System Information, System Restore, Task Scheduler, as well as all other tools covered in this chapter. Know the various disk management concepts.

Know the concepts related to disk management. You should know the various concepts related to disk management to initialize and mount a drive to assign a drive letter; the concepts of extending, splitting, and shrinking partitions (volumes); fault tolerance arrays; and Storage Spaces. GPT should be used with UEFI/BIOS, addressing 2 TB and larger drives, or if more than 26 logical volumes are required.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which MMC snap-in allows you to start a program based on a condition?
 - A. Programs and Features
 - B. Task Scheduler
 - C. Disk Management
 - D. Group Policy Editor
2. You need to select a format for a flash drive that is compatible with other operating systems. Which of the following is the filesystem should you choose?
 - A. NTFS
 - B. EFS
 - C. FAT32
 - D. CDFS
3. Which utility can be used to schedule a remote shutdown?
 - A. taskmgr
 - B. kill
 - C. shutdown
 - D. netstat
4. Which event log should be checked if you suspect that the computer has rebooted unexpectedly?
 - A. Application
 - B. Boot
 - C. Security
 - D. System
5. Which of the following filesystems is a proprietary filesystem created by Microsoft for use with large flash drives?
 - A. GPT
 - B. NFS
 - C. ext3
 - D. exFAT
6. Which section in the Settings app lets you configure backups?
 - A. System
 - B. Devices
 - C. Apps
 - D. Accounts

7. Which of the following filesystems is a proprietary filesystem created by Apple that offers SSD support?
 - A. HFS
 - B. ext4
 - C. APFS
 - D. exFAT
8. Which Control Panel applet would you use to join the operating system to a domain?
 - A. System
 - B. User Accounts
 - C. Windows Tools
 - D. Programs And Features
9. What can you do if a program is not responding to any commands and appears to be locked up?
 - A. Open the System Control Panel applet and choose Performance to see what process is causing the problem.
 - B. Add more memory.
 - C. Press Ctrl+Alt+Delete to reboot the computer.
 - D. Open Task Manager, select the appropriate task, and then click End Task.
10. Which command will start the Event Viewer snap-in?
 - A. `eventviewer.exe`
 - B. `eventvwr.msc`
 - C. `lusrmgr.msc`
 - D. `devmgmt.msc`
11. You need to configure an iSCSI connection; where will you find iSCSI settings?
 - A. Storage Spaces applet
 - B. Disk Management snap-in
 - C. Windows Tools
 - D. Device Manager snap-in
12. Which type of resources do you configure in Device Manager?
 - A. Hardware
 - B. Files and folders
 - C. Applications
 - D. Memory

13. What should be periodically done with mechanical hard drives to speed up performance when launching and saving files?
 - A. Free up space
 - B. Empty Recycle Bin
 - C. Defragment
 - D. Trim
14. You have a number of files with a .des extension and you want to be able to search each file's metadata; which applet should this be configured in?
 - A. Indexing Options applet
 - B. File Explorer Options applet
 - C. Windows Tools applet
 - D. Ease of Access Center applet
15. Which of the following partitions is specifically the partition from which the operating system boots?
 - A. Primary partition
 - B. Extended partition
 - C. Dynamic partition
 - D. Logical partition
16. Which of the following Registry hives contains information about the computer's hardware?
 - A. HKEY_CURRENT_MACHINE
 - B. HKEY_LOCAL_MACHINE
 - C. HKEY_MACHINE
 - D. HKEY_RESOURCES
17. Which of the following replaces the MBR in 64-bit versions of the Windows operating system?
 - A. LILO
 - B. NFS
 - C. GPT
 - D. GRUB
18. Which of the following utilities will rearrange the files on your hard disk to occupy contiguous chunks of space?
 - A. Disk Defragmenter
 - B. Windows File Explorer
 - C. Scandisk
 - D. Windows Backup

19. Which power mode will write the contents of the RAM to the disk?
- A. Hibernate
 - B. Standby
 - C. Sleep
 - D. Suspend
20. Which partition style allows up to 2,000 volumes to be partitioned on a physical disk?
- A. GPT
 - B. MBR
 - C. Dynamic Disk
 - D. Basic Disk

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

You are working at a company that has standardized on Windows 10 workstations for all. The phone rings, and it is your supervisor. He tells you that his workstation is running incredibly slowly, almost to the point where it is unusable. When you ask what he is running, he reports that he has exited out of everything but the operating system. You suspect there are background processes tying up the CPU and memory. Which utility can you have him use to look for such culprits?

Chapter 3

Windows Administration

**THE FOLLOWING COMPTIA A+ 220-1202
EXAM OBJECTIVES ARE COVERED IN THIS
CHAPTER:**

- ✓ 1.0 Operating Systems
- ✓ 1.2 Given a scenario, perform OS installations and upgrades in a diverse environment.
 - Boot methods
 - Universal Serial Bus (USB)
 - Network
 - Solid-state/flash drives
 - Internet-based
 - External/hot-swappable drive
 - Internal hard drive (partition)
 - Multiboot
 - Types of installations
 - Clean install
 - Upgrade
 - Image deployment
 - Remote network installation
 - Zero-touch deployment
 - Recovery partition
 - Repair installation
 - Other considerations
 - Third-party drivers



✓ **1.5 Given a scenario, use the appropriate Microsoft command-line tools.**

- Navigation
 - cd
 - dir
- Network
 - ipconfig
 - ping
 - netstat
 - nslookup
 - net use
 - tracert
 - pathping
- Disk Management
 - chkdsk
 - format
 - diskpart
- File Management
 - md
 - rmdir
 - robocopy
- Informational
 - hostname
 - net user
 - winver
 - whoami
 - [command name] /?
- OS Management
 - gpupdate
 - gpresult
 - sfc



✓ **1.7 Given a scenario, configure Microsoft Windows networking features on a client/desktop.**

- Domain joined vs. workgroup
 - Shared resources
 - Printers
 - File servers
 - Mapped drives
- Local OS firewall settings
 - Application restrictions and exceptions
 - Configuration
- Client network configuration
 - Internet Protocol (IP) addressing scheme
 - Domain Name System (DNS) settings
 - Subnet mask
 - Gateway
 - Static vs. dynamic
- Establish network connections
 - Virtual private network (VPN)
 - Wireless
 - Wired
 - Wireless wide area network (WWAN)/cellular network
- Proxy settings
- Public network vs. private network
- File Explorer navigation—network paths
- Metered connections and limitations



The previous chapter introduced the basic components of Windows operating systems and discussed the various tools used to configure Windows. This chapter builds on the previous chapters and focuses on Windows administration. In this chapter you will learn how to install and upgrade Windows 11 and learn various command-line tools. We'll also explore advanced concepts, such as Microsoft networking and design. All the content is generic to the Windows operating systems you'll be tested on during the 220–1202 exam. The content can be applied to any of the previous Windows operating systems.

Installing and Upgrading Windows



Exam Objective 1.2, “Given a scenario, perform OS installations and upgrades in a diverse environment,” requires you to know boot methods and the various types of installations. Be sure to know the various boot methods and the various types of installations for Windows 11.

Windows 11, released on October 5th, 2021, was the successor to Windows 10. The user interface (UI) looked familiar, and it was a well-needed upgrade. Microsoft also made the upgrade notification so irritating that many were forced to upgrade just to escape the notifications.

Regardless of the motivation to upgrade to Windows 11, its adoption was successful because of its meager system requirements. The hardware requirements for Windows 11 are similar to prior versions of Windows. However, starting with Windows 11, requirements have been scaled up for performance, as shown in Table 3.1. The meager hardware requirements were one of the driving factors for many to just hit the upgrade button.

TABLE 3.1 Windows 10/11 system requirements

Component	Windows 10 requirements	Windows 11 requirements
Processor	1 GHz or faster	1 GHz or faster (2 or more cores) 64-bit or System on a Chip (SoC)
RAM	1 GB (32-bit) or 2 GB (64-bit)	4 GB (64-bit)
Available Hard drive space	16 GB (32-bit) or 32 GB (64-bit)	64 GB (64-bit)

Component	Windows 10 requirements	Windows 11 requirements
Graphics card	DirectX 9 with WDDM 1.0 or higher driver	DirectX 12 with WDDM 2.0 driver
Display	800 x 600	High definition 720p
Additional hardware	N/A	UEFI & TPM 2.0 required

If you are planning to install Windows 11 on a computer, it must meet the specifications in Table 3.1, which are pretty easy to meet or exceed. However, sometimes you'll find that your organization's computers are older than you think. Windows 11 requires a realistic amount of computing power, RAM, and storage. In addition to these requirements, a UEFI firmware that is Secure Boot capable and a TPM 2.0 is required by Windows 11. To learn more about Windows 11 system requirements, visit <https://learn.microsoft.com/en-us/windows-hardware/design/minimum/minimum-hardware-requirements-overview>.


Checking Hardware Compatibility

In the past, Microsoft published a hardware compatibility list (HCL) for Windows XP and prior Windows operating systems. Microsoft went away from a published list of vendors in a static document and moved to an online version, which listed not only compatible hardware but compatible applications as well. You can access the Windows Compatible Products List at <https://partner.microsoft.com/en-us/dashboard/hardware/search/cpl>. However, the best way to check is to verify with the manufacturer that drivers exist and that the hardware is compatible with Windows 11.

Sometimes you need to deploy Windows 11 to a group of computers. Several tools can be used to collect information about the hardware of the current operating system. The easiest tool to use is the System Information utility. To access this utility, you simply log in and select System Information from the Start menu. You can use System Information to investigate the processor, RAM, hard drive, and video card, as well as specifics about each of the peripherals connected (see Figure 3.1). You can also access System Information by pressing Windows key+R, typing **msinfo32.exe**, and clicking OK.

You can also use System Information remotely. Select View and then choose Remote Computer from the drop-down menu and select the remote computer. You can also access this utility from the command line by typing **msinfo32.exe**. The **msinfo32.exe** command allows remote collection of information to a flat text file using the following command:

```
msinfo32.exe /computer computername /report c:\report.txt
```

FIGURE 3.1 System Information

The `msinfo32.exe` command can also be used to output the information to its native format using the following command:

```
msinfo32.exe /computer computername /nfo c:\report.nfo
```

However, this NFO file will need to be opened with the System Information utility, which might be difficult if there are more than 10 or so machines.

An alternative to the System Information (`msinfo32.exe`) utility is the Microsoft Assessment and Planning Toolkit (MAP), which can be downloaded from the Microsoft Download Center (www.microsoft.com/downloads). The MAP Toolkit allows for the automated inventory collection of hardware and software from the current operating systems. The MAP Toolkit can then produce a report from the collection of data and provide the administrator with the readiness of current hardware for Windows 11. The MAP Toolkit

requires the installation of a SQL database, which is included in the installation of the MAP Toolkit. The toolkit requires a dual-core 1.5 GHz processor, 2 GB of RAM, and a minimum of 1 GB of free hard drive space.


Windows Installation Options

Windows 11 can be installed as an upgrade or as a clean installation. When you choose Custom on Windows 11 23H2 or older, you can decide whether or not to format the hard disk. If you choose not to format the hard disk, the old operating system is placed in a folder called `WINDOWS.OLD`. When you choose to format the hard disk, it will erase your files, programs, and settings. You can also upgrade Windows 11 inside the existing operating system; we will cover how later in this chapter.

When installing Windows 10, you have the option to install it on *Basic Input/Output System (BIOS)*-based hardware or *Unified Extensible Firmware Interface (UEFI)*-based hardware. When installing Windows 11, you must install it on UEFI-based hardware. The hardware must support the newer standard of UEFI to install it in this fashion. UEFI hardware provides a feature called *Secure Boot*. Secure Boot operates by checking the signatures of the hardware, including the UEFI drivers (also called option ROMs), EFI applications, and, finally, the operating system. If the signatures are verified, the operating system is then given control of the boot process and the hardware. Windows 11 also requires that the hardware be Secure Boot-capable and that a TPM 2.0 module is installed.

For a UEFI installation, the partitioning of the drive will be laid out like Figure 3.2. The Recovery partition holds a bootable copy of the Windows Recovery Environment (WinRE) and is roughly 640 MB in size. The *EFI System Partition (ESP)* is a System partition used to hold the *Boot Configuration Data (BCD)* for the booting of the boot partition containing the *Windows kernel*.

FIGURE 3.2 Windows default disk layout





The Windows Recovery Environment (WinRE) is a bootable environment used by Windows for recovery of the operating system. WinRE is automatically booted when the operating system does not start properly after two consecutive boots. WinRE can also be booted into by restarting the operating system and holding down the Shift key. However, this process assumes that you can successfully boot the operating system.

The Installation Process

The elements of the installation of Windows 11 are almost identical to the installation of Windows 10 and prior versions of Windows. For that matter, it is similar to most operating systems, such as macOS or Linux. There are several common elements during the setup process that must be addressed, such as *locale* and where to install the operating system.


The installation of Windows 11 can be performed from a Windows 11 installation DVD-ROM. However, optical media such as DVD-ROM is rarely used because most laptops and tower computers no longer include optical drives. Universal Serial Bus (USB) installation media is a preferred method with most computer vendors. Installation media can be created with the Windows 11 Media Creation Tool, which can be downloaded from www.microsoft.com/en-us/software-download/windows11.



Windows XP was the last operating system to support CD-ROM installation media. Starting with Windows Vista, the size of every Windows release has grown and required DVD media for the installation. In addition to using the media creation tool to create USB media, you can download ISO files. The files can then be burned to DVD or used with a virtual machine by mounting the ISO file.



The version of Windows 11 as of the writing of this book is 24H2. Prior to Windows 11 24H2, the installation media would provide the means to install or upgrade when the media was booted. As of Windows 11 24H2, if the media is booted you only have the option to perform a clean installation of Windows 11. You can still upgrade with the installation media, but you must start the upgrade from within the currently installed operating system. We will cover the upgrade process later in this chapter. The following example explains a clean installation of Windows 11 with the latest Windows 11 24H2 media:

1. Start by booting the computer and directing it to boot from the physical media. Each computer will have a different way of entering the one-time boot menu, such as pressing the F12 key or the Esc key, or sometimes by initiating it from the BIOS/UEFI.
2. After the media is selected and detected, as the media boots you may see the message “Press any key to boot from CD or DVD.” This prevents the operating system from automatically booting into the installer. Once it boots, you will see the initial setup screen shown in Figure 3.3.

FIGURE 3.3 Select Language Settings screen


3. Select the language for the installation, as well as the time and currency format; these are often referred to as the *locale*. Click Next to continue.
4. On the next screen you will need to select the keyboard or input method language, as shown in Figure 3.4. Click Next to continue.
5. The next screen presents you with three options, as shown in Figure 3.5:
 - The first option is the most obvious: to install the operating system by selecting Install Windows 11.
 - The second option is to enter into recovery mode by selecting Repair Your PC.
 - The third option is less obvious: to exit the installation by pressing the X in the upper-right corner of the dialog box.

Before you can click the Next button, you must check the option “I agree everything will be deleted including files, apps, and settings.” Once you confirm by selecting the option, you can click Next.

FIGURE 3.4 Select Keyboard Settings screen**FIGURE 3.5** Select Setup Option screen


6. You will be prompted to enter the product key for the installation, as shown in Figure 3.6. You can skip the product key entry by selecting “I don’t have a product key.” This will allow you to bypass the product key screen and continue with the install. You will need to activate the operating system after it is installed.

FIGURE 3.6 Product Key screen



You can install Windows 11 without the product key. However, you will need to activate Windows 11 with a valid 25-digit product key to allow personal customization, such as wallpaper and Start menu settings. Windows 11 installs into *reduced functionality mode (RFM)*. If you don’t enter a product key, RFM will eventually nag you by displaying an Activate Windows message in the lower-right corner of the Desktop.

7. On the next screen, you will be asked which version of the operating system you want to install, as shown in Figure 3.7. Select to install only the operating system for which you have a valid license key for activation, since you cannot change editions later without a complete reinstallation of the operating system.



FIGURE 3.7 Windows edition selection

8. After selecting the edition, you will be prompted with the *end-user license agreement (EULA)*, as shown in Figure 3.8. To continue, read the license terms and if you agree, click Accept.



In previous versions of Windows 10 and prior to Windows 11 24H2, the next screen would prompt for the type of installation. The two options were Custom: Install Windows Only (Advanced) or Upgrade: Install Windows And Keep Files, Settings, And Applications. If you chose Custom, you could perform a clean installation of Windows, meaning that the drive would be wiped clean. Always remember that *custom* means a clean installation with Windows 11 23H2 and prior.

9. The next screen requires you to specify where you want to install the new operating system. You can delete, create, and extend partitions. If a partition exists and you are committed to a clean installation, then deleting the existing partitions is the most common task. In rare instances where no drives show up, you may need to install custom drivers specified by the vendor. You can perform that task from this screen as well by selecting Load Driver, as shown in Figure 3.9. Select the drive for the installation of Windows 11, and then click Next.


FIGURE 3.8 Windows end-user license agreement**FIGURE 3.9** Windows installation partitioning



The changes made in the partition utility are immediate.


10. You will be prompted with a Ready To Install confirmation screen. This screen will confirm that if you proceed you will install Windows 11 with the selected edition. The screen will also confirm that you will keep nothing, meaning it is a clean installation, as shown in Figure 3.10.

FIGURE 3.10 Ready To Install screen



11. This step is where the filesystem is formatted, the boot files are copied, and the operating system's files are applied to the disk. You will only see a progress counter, as shown in Figure 3.11.


When this process completes, the files for the operating system are on the hard drive but the operating system is not generalized to your computer; that happens in the next boot. After the computer reboots, it will detect hardware and run through what is called a *generalize pass*. The screen will show an installation progress percentage, as shown in Figure 3.12. During this process the various device drivers are being detected. This is the point where the operating system is adjusting itself to the computer hardware.

FIGURE 3.11 Windows installation progress**FIGURE 3.12** Windows Installing screen

Once the drivers are detected, the operating system will reboot again. During this boot, the drivers detected in the prior stage will be instantiated and the *specialize pass* will begin. In this pass, the locale (region) of the operating system will be chosen, as well as the keyboard layout. You can even choose a second keyboard layout, if you're switching out a keyboard.

12. Confirm the region by clicking Yes, as shown in Figure 3.13.

FIGURE 3.13 Operating system locale (region) setting



13. Confirm the keyboard layout by clicking Yes, as shown in Figure 3.14. It will also ask you if you want to add a second keyboard layout; just click Skip.

The Windows installation will display a message that it is checking for updates after the selection of the locale, keyboard layout, and additional keyboard layout question. The operating system will do some background work. The screen will assure you that something is going on in the background, and the message will change periodically. Then the operating system will reboot again.




NOTE

During the check for updates, the operating system will make a query to Microsoft Intune to check if the operating system's unique hash is contained in its global catalog. If the hash is in the global catalog, then

the respective organization will be queried for the Autopilot configuration. Windows Autopilot is a suite of cloud-based technologies designed to simplify the deployment and management of new Windows devices. This process eliminates the need for custom images and extensive setup, reducing the time and infrastructure required for device management. For more detailed information on Autopilot please refer to:

<https://learn.microsoft.com/en-us/autopilot/overview>

FIGURE 3.14 Operating system keyboard layout setting



14. When Windows 11 boots up, you will be given an opportunity to name your device during the setup, as shown in Figure 3.15. This is different from older operating systems, such as Windows 10, where a random name was created for you.

The operating system will seem to reboot after you confirm the name of the system. You will then see a Just A Moment screen while the operating system boots, as shown in Figure 3.16.

15. The next set of screens will help you set up the first user account as well as assign the administrator of the operating system.

FIGURE 3.15 Operating system name






FIGURE 3.16 Just A Moment screen




This is considered the OOBЕ pass, or the *out-of-box experience pass*, where the first account is set up for the operating system. You have the choice of logging in with a Microsoft account by choosing Set Up For Personal Use. This is great if you have the account combined with a Microsoft 365 subscription, since the licensing of the Office products are all streamlined when performing a login of this nature. However, if you want to join the computer to an organization's Intune mobile device management (MDM) service, you can select Set Up For Work Or School, as shown in Figure 3.17.

FIGURE 3.17 Windows account options



If you click Set Up For Personal Use, Windows 11 will start downloading Windows Updates and installing the updates, as shown in Figure 3.18. This ensures that your operating system is up to date when you first log in. Once updates are installed, you will be asked to sign in with a Microsoft account so that your apps, files, and services will sync to the device.

If you choose Set Up For Work Or School, you'll be asked to sign in with a work account or school account that is attached to an Intune service. This will give control of the device over to the organization, where it will be controlled by the MDM policies in Intune. Windows updates will not automatically install like the Set Up For Personal Use option, since the update policy will be decided for you by your organization.

FIGURE 3.18 Windows Updates installation



By choosing the Set Up For Work Or School option, you can choose to join the operating system to a domain. This will allow you to create a local admin account on the operating system.

We will install using a Microsoft account, so select the Set Up For Personal Use option, and click Next.

16. When the operating system is done installing Windows Updates, you will be presented with the Let's Add Your Microsoft Account dialog box, as shown in Figure 3.19. Enter your Microsoft account and password in the screens that follow, and select Next.
17. On the following screen you'll be asked to set a personal identification number (PIN). This allows you to easily log into the operating system. Click Create Pin. On the following screen create a new PIN, then type it again to confirm it, as shown in Figure 3.20, and click OK.




All new sign-in requests to the operating system can be performed with your PIN.

FIGURE 3.19 Microsoft account options for personal use**FIGURE 3.20** Set Up A PIN

18. The next screen asks you to Choose Privacy Settings For Your Device. This screen is known as *telemetry* options, as shown in Figure 3.21. Choose the level of privacy on the device as well as features like Location, Find My Device, Diagnostic Data, and other features that transmit to Microsoft's cloud. Once you have made your selection, click Accept to continue.



FIGURE 3.21 Choose Privacy Settings For Your device



19. If you have previously logged into another operating system with the Microsoft login, you will be presented with a Welcome Back screen, as shown in Figure 3.22. You can choose to Restore From This PC, or you can select More Options, then select Set Up As A New PC.


There are a number of advertisements to connect your phone, Microsoft 365 offers, OneDrive cloud storage offers, as well as informational advertisements. You can choose to click Decline to process or Continue to learn more. In addition, there are ads for Game Pass that you can choose to Skip For Now or Join. These ads will rotate as new versions of Windows are released, but the Microsoft 365, OneDrive, and Game Pass subscriptions will most likely remain.

20. Windows will display the Getting Things Ready For You screen while it configures the operating system. You will then be prompted with the Let's Customize Your Experience screen, as shown in Figure 3.23. You can choose how you will use the operating system, and a tailored experience will be provided. For example, you can choose development, gaming, entertainment, just to name a few. Choosing these will deliver ads, content, and tips to the operating system. You can always switch at a later time by selecting Start menu > Settings > Personalization > Device Usage. You can choose Skip to continue.

FIGURE 3.22 Restore From This PC screen**FIGURE 3.23** Let's Customize Your Experience

21. The next screen will ask you to opt into synchronizing your browsing data. The Always Have Access To Your Recent Browsing Data screen will allow you to opt into Microsoft Edge synchronization, as shown in Figure 3.24. Click Accept to continue.

FIGURE 3.24 Browsing data synchronization opt-in




22. The operating system will briefly state Finishing Up. Then you will be logged into the operating system. You will be greeted with a web page in Edge welcoming you to Windows, as shown in Figure 3.25.

Post-Installation Setup

After installation of the operating system, you have a few post-installation tasks. In this section we will cover all of the post-installation tasks that you should complete.

Time Zone Adjustment

At this point, you are pretty much done. The next screen you will see is the Windows 11 Desktop. However, throughout the entire setup, you were never asked for your time zone. The time zone is automatically calculated based on your IP address. If it is incorrect, then the following procedure will manually adjust it.


FIGURE 3.25 Welcome To Windows

To manually adjust the time, perform the following steps:

1. Right-click the clock in the lower-right corner of the Desktop.
2. Select Adjust Date And Time.
This will open the Date & Time screen, as shown in Figure 3.26.
3. Turn off Set Time Automatically if you want to change the time. If you want to change the time zone, turn off Set Time Zone Automatically if it is turned on.

Sign In To OneDrive


One of the key features of using a Microsoft account is the ability to back up your files with OneDrive. Before you can use this feature, you will need to log in. You can begin this process by clicking the gray cloud icon with the line through it, located in the system tray, as shown in Figure 3.27.

FIGURE 3.26 Adjusting Time Zone


1. Begin by clicking Sign In.
2. In the Set Up OneDrive dialog box, enter your email address that was set up with a Microsoft account. Then click Sign In.
3. Acknowledge the Privacy assurance screen by clicking Next.
4. Choose to send optional diagnostic usage data, then click Accept.
5. Confirm or change your OneDrive folder location, then click Next.
6. Confirm which folders OneDrive will back up, as shown in Figure 3.28. Then click Start Backup.
7. Choose to Upgrade your account for more storage, or select Not Now.
8. You can click Next through the tutorial, there are a few screens you will be presented with.

9. Click Later on the offer to Get The Mobile App.
10. Click Open My OneDrive Folder to exit the sign-in process.

FIGURE 3.27 Sign In To OneDrive



You will now notice that the line through the cloud has disappeared. Keep in mind that you will be limited to the free 5 GBs of space all Microsoft accounts are allotted. If you need more storage, you will need to purchase a Microsoft 365 account or OneDrive storage. More information can be found at www.microsoft.com/en-us/microsoft-365/onedrive/compare-onedrive-plans.

FIGURE 3.28 Back Up Folders On This PC

Sign In To Edge

When you launch Edge for the first time you will be asked to sign in. This allows you to sync favorites and browsing data. You can begin the process by clicking Sign In To Sync Data in the upper-left corner of the Edge browser, as shown in Figure 3.29.



If you have already signed in to OneDrive, the operating system will detect your sign-in and sign you into Edge automatically.

The Upgrade Process

Now let's take a look at the in-place upgrade process for a Windows 10 operating system to Windows 11. It is recommended that you start by performing all of the current Windows Updates first. During Windows Update, you may see an ad to upgrade to Windows 11, as shown in Figure 3.30. You can begin the upgrade from this ad by selecting Download And Install.

The tried and true method of upgrading Windows is by inserting the Windows 11 media or connecting to a network share that contains the media and launching `setup.exe`. This will require answering a User Account Control (UAC) prompt. When the setup process starts, it will give you the option to change how Windows Setup downloads updates, as shown in Figure 3.31, or you can just click Next at this point. The default is to proceed with the download of Windows Updates for the installation.




FIGURE 3.29 Signing in to Edge**FIGURE 3.30** Download And Install Windows 11

FIGURE 3.31 Install Windows 11 screen

Before proceeding, you must accept the end-user license agreement (EULA), also known as the license terms, as shown in Figure 3.32.

The installer will check and download updates necessary for the installation before continuing, as shown in Figure 3.33. This will ensure that you are secure and the upgrade process is smooth and without complication.

The next screen is the Ready To Install screen, shown in Figure 3.34. This confirms the edition detected and also confirms that personal files and apps will be kept. You have the option to change what is kept during the upgrade process by clicking Change What To Keep. You can choose one of these options: Keep Personal Files And Apps, Keep Personal Files Only, or Nothing.

The upgrade process will begin, and you will see the familiar progress percentage in the upper-left corner of the screen, as shown in Figure 3.35. The computer will reboot, and Setup will continue with a different progress screen, as shown in Figure 3.36. The computer will reboot several times during this process. As mentioned earlier, device drivers will be detected, and a reboot is required for the drivers to be properly loaded and instantiated.



FIGURE 3.32 Windows 11 end-user license agreement**FIGURE 3.33** Windows 11 updates check

FIGURE 3.34 Windows 11 ReadyTo Install screen





FIGURE 3.35 Installing Windows 11 upgrade percentage



FIGURE 3.36 Windows 11 upgrade percentage after reboot

The Windows 11 upgrade process has omitted the possibility of choosing the wrong edition of Windows 11 by automatically detecting the currently installed edition. The Windows 11 upgrade process has also omitted the telemetry settings screen, inheriting the settings from the original install.

You have 10 days to revert to the prior versions of Windows after the upgrade. The previous operating system files are placed in the `C:\WINDOWS.OLD` folder. To roll back, click Start menu > Settings > System > Recovery > Go Back, and you can revert back to Windows 10.

Repair Installation

A *repair installation* is used when you want to reinstall the operating system without losing personal data files, application settings, or applications you've installed. The installation is similar to an upgrade, as described in the previous section, except that Windows 11 will detect that it is installed already. You will be presented with the default option of Keep Personal Files And Apps. However, you can choose one of these options: Keep Personal Files And Apps, Keep Personal Files Only, or Nothing. The setup process will then reinstall the OS without affecting your personal files, applications, and their corresponding settings. It will, however, reinstall the operating system files, so it is considered a *repair* installation.

Another option for reinstalling Windows 11 is to reset the operating system with the Reset This PC option. This option is used to reset the operating system back to its original state. It provides another way to fix the operating system when it looks to be corrupted. This method should be used as a last resort. It does allow you the choice to keep personal files or completely erase them along with the operating system. On Windows 11, you can reset the operating system with the Reset This PC option by clicking Start menu > Settings > System > Recovery, and then choosing Reset This PC.

Recovery Partition

Some vendors will supply a recovery partition that contains the original image for the OS the system came installed with. In the event the operating system is corrupted, the recovery partition can be used to reimage the device back to the factory image. When vendors started to include the recovery partition, the image could be restored by booting the computer into the recovery partition, and a third-party utility would be used to reimage the device. However, with Windows 11 the use of third-party utilities is no longer needed—Windows has a provision for locally installing the image via WinRE.

Side-by-Side Upgrading Windows 10/11

Side-by-side upgrades require additional hardware, since the original operating system will not be modified during the upgrade. The new hardware will be the hardware that is installed from scratch with Windows 11. The user settings and data will then be migrated over from the previous operating system to the new one.

Side-by-side upgrades are the best way to upgrade when you need to upgrade a system but the current system is in use. When dealing with one system, side-by-side upgrades are great because you don't have to worry about backups of the OS. However, applications must be reinstalled, so this approach does have its disadvantages. When multiple systems require upgrading, a *rolling upgrade* can be performed. A rolling upgrade is a variation on the side-by-side upgrade that allows the decommissioned equipment to be used to upgrade the next user, and this creates a continuous upgrade cycle. In other words, as you complete the upgrade of one user, their old system becomes the next user's new system, and the process continues until you get to the last user.

There are a few ways of migrating the user data from the old device to the new device, depending on how your deployment is configured. If the devices are using OneDrive to back up users' data folders, then users can simply sync their data onto their new devices. This method will work for data files and Microsoft Store apps, but OneDrive will not back up and restore Win32 applications.

If you are not using OneDrive, then you can migrate user data with Windows 11 using the Microsoft Windows User State Migration Tool (USMT). USMT allows you to migrate user file settings related to the applications, Desktop configuration, and accounts from one computer to another computer. The migration can be performed via a network connection or a hard drive.

USMT is compatible with Windows 10 and Windows 11, but it is currently available only as part of the Windows 11 Assessment and Deployment Kit (ADK). Although USMT is not part of the CompTIA objectives for the 220-1202 exam, it is important to know that a tool like USMT exists, because the migration of user information is a component of an upgrade.



Additional information about USMT can be found here:

<https://learn.microsoft.com/en-us/windows/deployment/usmt/usmt-overview>

Using USMT requires an investment of time. If you are performing only a few side-by-side upgrades, you can copy the user's profile directory manually. Then, when the new computer is ready, you can copy over the various folders, such as videos, pictures, and documents.

Image Deployment

Upgrading operating systems takes time and complex steps to complete, such as running Windows Upgrade, installing updates, upgrading the OS, and installing programs. These steps can take a tremendous amount of time, but if you have fewer than a dozen devices to upgrade, then you could manually upgrade the devices. However, if you have more than a dozen, there are alternatives such as image deployment.

The image deployment process is simple in concept. It starts with a base operating system on a reference computer utilizing real hardware or a virtual machine. The software is then installed on the operating system, an answer file is created using the Windows 11 Assessment and Deployment Kit (ADK), and the Sysprep tool is run on the operating system. The last step is to use a tool like Windows Deployment Services (WDS) to create an image of the operating system. Once the image is captured, it can be deployed to the other computers using the Windows Preinstallation Environment (WinPE).

Zero-Touch Deployment

Zero-touch deployments represent a modern approach to device management, automating the setup and configuration of new devices with minimal manual intervention. Two popular tools for zero-touch deployments are Microsoft Deployment Toolkit (MDT) and Windows Autopilot.

On-Premises Deployments

MDT is an on-premises solution that provides extensive customization and control over the deployment process, making it ideal for organizations with specific and complex deployment needs. Zero-touch deployments using MDT are designed to automate the deployment of Windows operating systems with minimal manual intervention. This process is often integrated with Microsoft Endpoint Configuration Manager (MECM) to enhance its capabilities. In a zero-touch deployment, MDT and MECM work together

to create a seamless deployment experience. The process begins with the creation of a deployment network share inside MDT, where all necessary files, applications, and drivers are stored. These resources are then packaged into a task sequence, which is a series of steps that automate the installation and configuration of the operating system on target devices.

One of the key components of zero-touch deployments is the use of task sequences. These sequences are highly customizable and can include steps such as partitioning the hard drive, applying the operating system image, installing applications, and configuring settings. The task sequence is executed automatically when the target device boots up, typically via a Preboot Execution Environment (PXE) boot or a bootable media. This automation ensures that each device is configured consistently according to the organization's standards, reducing the risk of human error and speeding up the deployment process.

Zero-touch deployments with MDT also support advanced features such as user state migration and application deployment. User State Migration Tool (USMT) can be integrated into the task sequence to capture and restore user data and settings, ensuring a smooth transition for end users. Additionally, MDT can deploy applications as part of the task sequence, either by installing them directly from the deployment share or by downloading them from a network location. This flexibility allows IT administrators to tailor the deployment process to meet the specific needs of their organization, making zero-touch deployments a powerful and efficient solution for managing large-scale Windows deployments.

Cloud-Based Deployments

Zero-touch deployments of Windows 11 leverage cloud-based technologies to streamline the setup and configuration of new devices, minimizing the need for IT intervention. This approach is primarily facilitated by Windows Autopilot, a feature within Microsoft Intune, which allows organizations to preconfigure devices before they even reach the end user. When a user receives their device, they simply power it on and connect to the Internet, and the device automatically downloads and applies the necessary configurations, applications, and policies. This ensures a consistent and secure setup for all devices.

One of the key benefits of zero-touch deployments is the enhanced user experience. Employees can start using their new devices almost immediately, without waiting for manual setup or configuration. This is particularly advantageous for remote or hybrid work environments, where employees may not have easy access to IT support. Additionally, zero-touch deployments help maintain security and compliance by ensuring that all devices are configured according to the organization's policies from the day they are setup. This reduces the risk of misconfigurations and potential security vulnerabilities.

Zero-touch deployments are highly scalable, making them ideal for organizations of all sizes. Whether deploying a handful of devices or thousands, the process remains efficient and manageable. IT administrators can create and manage deployment profiles, monitor the status of deployments, and make adjustments as needed, all from a

centralized console. This scalability, combined with the automation and consistency provided by Windows Autopilot, makes zero-touch deployments a powerful tool for modern IT management.



Deployment of the Windows operating system can be dedicated to a book itself. The CompTIA exam requires basic knowledge of zero-touch deployments. If you want to learn more about AutoPilot, visit <https://learn.microsoft.com/en-us/autopilot/overview>. To learn more about the MDT you can visit <https://learn.microsoft.com/en-us/mem/configmgr/mdt>. Information about MECM can be found at <https://learn.microsoft.com/en-us/mem/configmgr>.

Upgrading Editions of Windows

Windows 11 has the ability to upgrade at any time from one edition of the operating system to a higher one (e.g., from Windows 11 Home to Windows 11 Pro). This can easily be accomplished by entering the appropriate activation product key. You can access the activation screen by choosing Start menu > System > Activation and selecting Change from the Change Product Key menu. You can only upgrade editions; downgrading of retail editions is not supported. However, downgrading of volume license editions, such as Windows 11 Enterprise to Windows 11 Pro, can be achieved. Downgrading of Windows 11 Education to Windows 10/11 Pro can also be done. Downgrading of volume license editions is not formally supported by Microsoft but can be accomplished.

Security and Feature Updates

The Windows platform has always had security updates and feature updates. In older versions of Windows such as XP, the features were subtle additions to the operating system. However, with the introduction of Windows 10 entire releases have been dedicated to major feature releases, such as the Creators update and the Anniversary update. Since then, the features have not been so dramatic, but Windows continues to get feature updates with every major/minor update.

Windows versions typically change twice a year (semi-annually) and have done so since the introduction of Windows 10. They just haven't been obvious, because they were downloaded as a Windows Update. The original version of Windows 10 is 1507, which was released in July 2015. You may have noticed the pattern. The version is a *date code*, consisting of the last two digits of the year (15) and the two-digit month (07). So, it's simple to calculate when the last major update was released, and what is currently installed. With the October 2020 release of Windows 10, Microsoft deviated from this naming convention, using H1 for the first half and H2 for the second half of the year. For example, version 21H1 was released in the first half of 2021. Windows 11 also follows the same date code with its initial release in October 2021; its date code is 21H2.


Versions are updated twice a year (semi-annually), usually in spring and fall. They are often referred to as the Windows 11 Spring Update or Windows 11 Fall Update, respectively. In the past few years Microsoft has only released an H2 for Windows 11, which equates to one major update released per year, in the fall. As of this writing, the current Windows 10 version is 22H2, and the Windows 11 version is 24H2. They were both used in the development of this book. Using `winver.exe`, you can see the actual version of the operating system, as shown in Figure 3.37.

FIGURE 3.37 Discovering your version of Windows 11



The life cycle of a Windows product has an end date for support, generally a year to two years after the product has been released. Once the end date is met, Microsoft will no longer support that version of Windows. The retirement date of Windows is generally 10 to 12 years, depending on the operating system's popularity. When the product reaches its retirement date, security updates are no longer furnished for the operating system. Therefore, you should be planning OS upgrades on a consistent schedule.

Probably the biggest change to Windows 11 is the way updates are delivered to the operating system. The end user no longer has the option to choose if they want to set up Windows Updates; they are mandatory and unavoidable. You can pause updates for 1 week to 5 weeks (35 days), but inevitably updates will be installed. To pause Windows Updates, click the Start menu > Settings > Windows Update > Pause Updates. Then you can choose to pause updates for up to 5 weeks. Figure 3.38 shows the options to pause Windows Updates.

FIGURE 3.38 Pause options for Windows Updates

Windows 11 has three different build branches of code (updates) that will be installed on a regular basis with automatic updates. The branches are as follows:

General Availability Channel This is the default servicing channel for Windows 11 devices, with the exception of Windows 11 Enterprise LTSC (Long-Term Servicing Channel).

Insider Program The Insider Program allows you to get the latest product updates from Microsoft before they are available for the targeted release. You must enroll your organization first, and then you can enroll your device to get the latest updates from the Windows Insider Program build channel.

Long-Term Servicing Channel Only users of the Windows 11 Enterprise LTSC Edition can select this channel for updates. The operating system will receive monthly updates via Windows Updates, but no new features will be added over the lifetime of the product.

Table 3.2 details the list of the various Windows 11 editions and the options for changing their service update channels.

TABLE 3.2 Windows servicing channel options

Windows 11 Edition	General Availability Channel	Insider Program	Long-Term Servicing Channel
Home	Yes	Yes	No
Pro	Yes	Yes	No
Enterprise	Yes	Yes	No
Enterprise LTSC	No	No	Yes
Education	Yes	Yes	No



With the introduction of Windows 10 and the initial Windows Update servicing channel options, Microsoft has introduced the concept of Windows as a service (WaaS). A deeper dive on Windows servicing channels can be viewed here:

<https://learn.microsoft.com/en-us/windows/deployment/update/waas-overview>

Installation/Upgrade Boot Options

You can begin the installation or upgrade process by booting from a number of sources. There are three sources in particular with which you should be familiar: USB flash drive, CD-ROM/DVD ISO image, and network boot (PXE). The one method most commonly used for an attended installation is the USB flash drive boot. Another common installation method is booting from ISO image. This installation method is used with virtual machines, since the DVD-ROM is in the form of an ISO image and can be stored on a network filesystem.

You can boot a PC over the network (rather than from an ISO image, USB, or hard disk) with *Windows Preinstallation Environment (WinPE)*, which is a stub operating system that creates a *Preboot Execution Environment (PXE)*. A stub operating system is characterized as a scaled-down version of the primary operating system. A stub operating system usually has basic functionality to perform a basic task, such as the installation of the Windows OS. When Windows 11 is installed with this method, it is considered a *remote network installation*.

You can install or upgrade Windows with traditional installation media, such as DVD ISO image, USB flash drives, and network boot (PXE), but they all tend to be a bit slow. An alternate method of installation and upgrade of Windows is from an external/hot-swappable drive, such as a USB hard drive. You can even use an external hot-swappable drive, such as an eSATA. However, the top speed on eSATA is 6 Gbps, compared to the top speed of USB 3.1 at 10 Gbps and USB 4 at 80 Gbps!

There is one other final option for installing Windows that is often used for reinstallation of the operating system by the original equipment manufacturer (OEM). That option is an internal hidden partition that contains the installation media. The OEM will often include this partition on the hard drive to allow the user to easily reinstall the operating system in the event of problems. In recent years, OEMs have moved away from an internal hard drive partition in favor of the Reset Your PC option in Windows.

Installing or booting Windows 11 can be accomplished through several versatile methods, each catering to different needs and scenarios. Whether you're setting up a new system, upgrading an existing one, or managing multiple devices in an organization, understanding these methods can streamline the process and ensure a smooth installation. From traditional USB drives to advanced network booting and Internet-based installations, each method offers unique advantages. This chapter will walk you through the technical details of each approach, helping you choose the best method for your specific requirements. Let's explore the various ways to get Windows 11 up and running on your device.

Windows 11 Boot Process

The Windows 11 boot process is a sequence of steps that the operating system follows to initialize and prepare the system for use. This process begins when the computer is powered on and continues until the Windows Desktop is fully loaded and ready for user interaction. Understanding these steps can help diagnose and troubleshoot boot-related issues. The following is a typical boot process for Windows 11:

1. Power-On Self-Test (POST)

When the computer is powered on, the UEFI firmware performs a Power-On Self-Test (POST) to check the hardware components. This step ensures that all critical hardware, such as the CPU, memory, and storage devices, are functioning correctly.

2. UEFI Initialization

After a successful POST, the UEFI firmware looks for the boot loader in the EFI System Partition (ESP). The firmware executes the Windows Boot Manager (`bootmgfw.efi`) found in this partition.

3. Boot Configuration Data (BCD) Reading

The Windows Boot Manager reads the Boot Configuration Data (BCD) to determine which operating system to load. The BCD contains essential information about the operating system's location and boot parameters.

4. Loading the Windows Boot Loader

The Boot Manager loads the Windows Boot Loader (`winload.efi`), which is responsible for initializing the Windows kernel. This step sets the stage for the operating system to take control.

5. Kernel Initialization

The Windows Boot Loader initializes the kernel (`ntoskrnl.exe`), the core component of the operating system. The kernel then initializes the hardware abstraction layer (HAL), essential drivers, and the system registry.

6. Session Manager Subsystem (smss.exe)

Control is handed over to the Session Manager Subsystem (smss.exe), which starts the user session. The Session Manager initializes system processes, loads necessary drivers and services, and sets up the virtual memory paging file.

7. Winlogon Process (winlogon.exe)

The Winlogon process is started, handling user logins and launching the Windows shell. This process is crucial for managing user authentication and starting the graphical user interface.

8. Windows Shell Initialization

The Windows shell (explorer.exe) loads the Desktop environment, including the Taskbar, Start menu, and Desktop icons. This step prepares the system for user interaction.

9. User Interaction

At this point, the system is fully booted, and the user can begin interacting with Windows 11. The Desktop environment is ready, and all necessary background services are running.

Each of these steps is critical for a successful boot. Any failure along the way can result in boot errors or a system that fails to start properly. The installation process follows a similar boot process, the Windows Preinstallation Environment (WinPE) is a mini-Windows operating system used to install Windows 11 and prior versions of Windows. Understanding this process can help in diagnosing issues and ensuring a smooth boot experience.

Universal Serial Bus (USB)

Installing Windows 11 via a USB drive is a widely used and efficient method. To begin, you'll need a USB flash drive with at least 8 GB of storage. Download the Windows 11 ISO file from the official Microsoft website. Use a tool like Rufus to create a bootable USB drive by selecting the ISO file and the USB drive as the target. Ensure the partition scheme is set to GPT for UEFI systems or MBR for BIOS systems. Once the USB drive is ready, insert it into your computer, restart the system, and enter the BIOS/UEFI settings (usually by pressing F2, F12, Delete, or Esc during startup). Set the USB drive as the primary boot device. Save the changes and exit the BIOS/UEFI. Your computer will boot from the USB drive, and you can follow the onscreen instructions to complete the installation process.

Network

Network booting, or PXE (Preboot Execution Environment), allows you to install Windows 11 over a network connection, which is particularly useful for deploying the OS across multiple machines. PXE booting requires the installation and configuration of a DHCP server to provide network boot options and a TFTP server to host the Windows 11 installation files. On the client machines, enable PXE boot in the BIOS/UEFI settings. When the client

computers start, they will receive the network boot information from the DHCP server, connect to the TFTP server, and download the necessary boot files. The Windows Deployment Services (WDS) can be used to manage and automate this process, making it easier to deploy Windows 11 across an organization.

Solid-State/Flash Drives

Installing Windows 11 from a solid-state drive (SSD) or flash drive offers faster read/write speeds compared to traditional hard drives. To prepare the SSD or flash drive, format it using FAT32 and ensure it contains the Windows 11 installation files. You can use tools like Disk-Part to clean and format the drive. Connect the SSD or flash drive to your computer, restart the system, and access the BIOS/UEFI settings. Set the SSD or flash drive as the primary boot device. Save the changes and exit the BIOS/UEFI. Your computer will boot from the SSD or flash drive, and you can follow the installation prompts to complete the process, benefiting from quicker data transfer rates and reduced installation time.

Internet-Based

Internet-based installation of Windows 11 leverages cloud services to download and install the operating system directly from the Internet. This method is ideal for devices with limited local storage or for users who prefer not to use physical media. To perform an Internet-based installation, connect your device to a stable internet connection. Access the Windows 11 download page www.microsoft.com/en-us/software-download/windows11 and download the Windows Installation Assistant. Run the Windows 11 Installation Assistant, which will guide you through the process of downloading and installing Windows 11. This method ensures you have the latest updates and drivers during the installation process.

External/Hot-Swappable Drive

Using an external or hot-swappable drive to install Windows 11 provides flexibility and ease of use. These drives can be connected and disconnected without shutting down the system. To install Windows 11, connect the external drive to your computer and ensure it contains the installation files. Format the drive using NTFS if necessary. Access the BIOS/UEFI settings and set the external drive as the primary boot device. Save the changes and exit the BIOS/UEFI. Restart your computer, and it will boot from the external drive. Follow the on-screen instructions to complete the installation. This method is particularly useful for systems that require frequent OS changes or testing, such as a development or testing lab.

Internal Hard Drive (Partition)

Booting and installing Windows setup from an internal hard drive partition allows you to reimagine the computer rapidly. This is commonly known as a recovery partition, as previously discussed. To create your own recovery partition, you would use unpartitioned space on the drive to create a primary partition and set it to be an active partition. Setting it to active would allow it to be bootable in the BIOS. Next you will format the space with NTFS,

then copy the Windows 11 setup files to the new partition. The `bcdedit.exe` tool can then be used to register the new partition as a boot option.

This method was commonly used in prior operating systems, because Windows lacked the tools to recover the operating system. However, starting with Windows 10, a recovery partition is created, and this is no longer needed. You can still perform this type of install, but the built-in functionality works the same.

Multiboot

Multibooting allows you to install multiple operating systems on a single computer, giving you the option to choose which OS to boot into at startup. To set up a multiboot system with Windows 11, first ensure you have separate partitions for each operating system. Use disk management tools to create or resize partitions as needed. Install the first OS as usual, then install Windows 11 on a different partition. During the installation, the Windows boot manager will be configured to allow you to select the desired OS each time you start your computer. This setup is ideal for users who need to run different operating systems for various tasks, such as development, testing, or personal use.

Multiboot systems were a mainstream option 20 years ago to boot into two separate operating systems on one computer. However, with the introduction of virtualization, such as Hyper-V, multiboot operating systems have lost their utility. Starting with Windows Vista SP2 released in 2009, Hyper-V was a downloadable feature.

Considerations

With every great plan there are unplanned consequences; we call them *considerations*. When installing Windows 11 or upgrading to Windows 11, there are several considerations surrounding the applications you need to use and the hardware Windows 11 is being installed on.

When you upgrade an operating system, you have the potential for data loss. Therefore, it is always advisable to back up files and user preferences before starting an upgrade. Depending on the user you are upgrading in your organization, if files are lost it could cause catastrophic loss of sales, payments, and most importantly, time. Whenever possible, perform a full-drive backup, because settings may be in a spot you wouldn't normally back up. If a full-drive backup cannot be performed, then perhaps replace the drive with another one. Label the original drive *removed* and recycle it after you confirm that everything is acceptable with the end user.

The applications must be supported on Windows 11. You may think that if it runs, then it's supported, but you'd be wrong. Just because a program runs on Windows 11 doesn't mean that it was meant to be run on the operating system. It is up to the discretion of the software vendor if they will support you on the latest operating system. This is often discovered in the event you need help from them. So always check if the application is supported on the latest OS before upgrading. In many cases you'll find that software vendors require the latest version of their product purchased or installed for it to be supported on the latest

OS. Applications are also backward compatible with older operating systems, because not everyone will be on the latest and greatest operating system.

The hardware that Windows 11 is running on is another consideration. Many motherboards and peripherals need third-party drivers. These drivers must be supported on Windows 11 or they may not function correctly. Older network interface cards (NICs) are notorious for not being supported on Windows 11. So, always check the hardware vendor's website before upgrading to Windows 11.

Command-Line Tools

Although the exam focuses on the Windows operating systems, it tests a great number of concepts that carry over from the Microsoft Disk Operating System (MS-DOS). MS-DOS was never meant to be extremely user friendly. Its roots are in Control Program for Microcomputers (CP/M), which, in turn, has its roots in UNIX. Both of these older OSs are command line-based, and so is MS-DOS. In other words, they all use long strings of commands typed in at the computer keyboard to perform operations. Some people prefer this type of interaction with the computer, including many folks with technical backgrounds. Although Windows has left the full command-line interface behind, it still contains a bit of DOS, and you get to it through the command prompt.

Although you can't tell by looking at it, the Windows command prompt is actually a Windows program that is intentionally *designed* to have the look and feel of a DOS command line. Because it is, despite its appearance, a Windows program, the command prompt provides all the stability and configurability you expect from Windows. You can access a command prompt by running `cmd.exe`.

A number of diagnostic utilities are often run at the command prompt. They can be broken into two categories: networking and operating system. The utilities associated with networking appear in other chapters, but the focus here is on the utilities associated with the operating system.



Exam Objective 1.5, "Given a scenario, use the appropriate Microsoft command-line tools," requires you to know various command-line tools. The OS command-line tools that you are expected to know for the exam are `cd`, `dir`, `ipconfig`, `ping`, `netstat`, `nslookup`, `net use`, `tracert`, `pathping`, `chkdsk`, `format`, `diskpart`, `md`, `rmdir`, `robocopy`, `hostname`, `net user`, `winver`, `whoami`, `[command name] /?`, `gpupdate`, `gpresult`, and `sfc`.

In the following sections we will cover a number of commands you will see on the exam, along with the commands available with standard privileges, as opposed to those with administrative privileges.

Navigation and File Management

The power of the command line comes from the level of detail that can be attained with simple commands. There are several commands that allow you to extract a high level of detail. However, the command line is an unforgiving user interface that requires specific commands, whereas with the GUI you just need to know how to point and click, but it lacks the high level of detail. In the following section we will cover basic navigation at the command line and basic file management commands.

dir Command

The `dir` command is used to display a list of the files and folders/subdirectories within a directory. When you use it without any parameters, `dir` will show you not only that information but also the volume label and serial number, along with the amount of free space, in bytes, remaining on the disk.

Wildcards can be used with the command to list all files that begin with a certain letter or end with certain letters. An example is typing `dir *.txt` to list all the text files in a directory. A plethora of parameters are available that can be used to customize the results or the display. Table 3.3 lists some of the most common switches available for `dir`.

TABLE 3.3 Common `dir` switches

Switch	Purpose
/a	Allows you to specify the attributes of files you are seeking (hidden, system, and so on)
/o	Allows you to specify a different display order (alphabetic is the default)
/l	Returns the results unsorted and in lowercase format
/s	Recursively searches through subdirectories as well as the current directory
/t	Sorts the files according to time order
/p	Displays the results one page/screen at a time
/q	Shows file ownership

cd, md, rd Commands

The `cd`, `md`, and `rmdir` commands are used to change (or display), make, and remove directories, respectively. The commands `cd`, `md`, and `rd` are shorthand versions of the `chdir`, `mkdir`, and `rmdir` commands, respectively. Table 3.4 lists their usage and switches.

TABLE 3.4 cd/md/rmdir usage and switches

Command	Purpose
cd [path]	Changes to the specified directory.
cd /d [drive:][path]	Changes to the specified directory on the drive.
cd ..	Changes to the directory that is up one level.
cd \	Changes to the root directory of the drive.
md [drive:][path]	Makes a directory in the specified path. If you don't specify a path, the directory will be created in your current directory.
rmdir [drive:][path]	Removes (deletes) the specified directory.
rmdir /s [drive:][path]	Removes all directories and files in the specified directory, including the specified directory itself.
rmdir /q [drive:][path]	Quiet mode. You won't be asked whether you're sure you want to delete the specified directory when you use /s.

Drive Navigation

So far, you've seen the basics of looking at directories with the `dir` command, changing directories, making directories, and removing them. However, up to this point we have assumed you are on the same partition. The `cd` command will change directories within a drive letter, such as the C: drive, but it will not change drive letters unless you supply the `/d` switch. To change drives without using the `dir` command, just enter the drive letter and append a semicolon to it. For example, if you want to change to the D: drive, enter **d:** at the command prompt. You can then use the `cd` command followed by the drive letter, and if you want to change back, enter **c:** at the command prompt.

Now that you've learned how to navigate the command prompt to look at files, let's use that knowledge in Exercise 3.1.

EXERCISE 3.1

Command-Line Directory Management

1. Open a command prompt by selecting the Start menu, typing **cmd** in the search box, right-clicking the result Command Prompt (System), and selecting Run As Administrator. Alternately you can press the Windows key+R, type **cmd**, then press Ctrl+Shift+Enter to run as administrator.
2. Change to the root of your C: drive by typing **cd /d C:** and pressing Enter.

EXERCISE 3.1 (continued)

Note: If you are already in C:, all you have to do is type **cd ** and press Enter.

3. Create a directory called C16 by typing **md C16** and pressing Enter.
4. Change to the C16 directory by typing **cd C16** and pressing Enter.
5. Create several layers of subdirectories at once. Type **md A1\B2\C3\D4** and press Enter. Notice that these commands create each of the directories that you specified. You now have a directory structure that looks like this: C:\C16\A1\B2\C3\D4.
6. Change back to your root directory by typing **cd **.
7. Attempt to delete the C16 directory by typing **rd C16** and pressing Enter.

Windows won't let you delete the directory because the directory is not empty. This is a safety measure. Now let's really delete it.

8. Delete the C16 directory and all of its subdirectories by typing **rmdir /s C16** and pressing Enter.
9. You will be asked whether you're sure that you want to delete the directory. If you are, type **Y** and press Enter.
10. To close the command prompt window, type **exit**.

Note that if you had used the **/q** option in addition to **/s**, your system wouldn't have asked whether you were sure; it would have just deleted the directories.

copy Command

The copy command does what it says: It makes a copy of a file in a second location. (To copy a file and then remove it from its original location, use the move command.) Here's the syntax for copy:

```
copy [filename] [destination]
```

It's pretty straightforward. There are several switches for copy, but in practice they are rarely used. The three most commonly used switches are **/a**, which indicates an ASCII text file; **/v**, which verifies that the files are written correctly after the copy; and **/y**, which suppresses the prompt asking whether you're sure that you want to overwrite files if they exist in the destination directory.



The copy command cannot be used to copy directories. Use xcopy for that function.



One useful tip is to use wildcards. For example, at the command prompt, the asterisk (*) is a wildcard that means *everything*. So, you could type **copy *.exe** to copy all files that have an .exe filename extension, or you could type **copy *.*** to copy all files in your current directory. The other popular wild card is the question mark (?), which does not mean everything but instead means one thing. For example, copy **ABC?.exe** would copy only .exe files with four-letter names, of which the first three letters are ABC.

xcopy Command

If you are comfortable with the copy command, learning xcopy shouldn't pose too many problems. It's basically an extension of copy with one notable exception—it's designed to copy folders as well as files. The syntax is as follows:

```
xcopy [source] [destination] [switches]
```

There are 26 xcopy switches. Some commonly used ones are listed in Table 3.5.

TABLE 3.5 xcopy switches

Switch	Purpose
/a	Copies only files that have the Archive attribute set and does not clear the attribute (useful for making a quick backup of files while not disrupting a normal backup routine)
/e	Copies directories and subdirectories, including empty directories
/f	Displays full source and destination filenames when copying
/g	Allows copying of encrypted files to a destination that does not support encryption
/h	Copies hidden and system files as well
/k	Copies attributes (by default, xcopy resets the Read-Only attribute)
/o	Copies file ownership and ACL information (NTFS permissions)
/r	Overwrites read-only files
/s	Copies directories and subdirectories but not empty directories
/u	Copies only files that already exist in the destination
/v	Verifies the size of each new file

Perhaps the most important switch is `/o`. If you use `xcopy` to copy files from one location to another, the filesystem creates new versions of the files in the new location without changing the old files. In NTFS, when a new file is created, it inherits permissions from its new parent directory. This could cause problems if you copy files. (Users who didn't have access to the file before might have access now.) If you want to retain the original permissions, use `xcopy /o`.

robocopy Command

`robocopy.exe`, short for Robust File Copy, is a powerful command-line utility included with recent versions of Windows designed to handle complex file-copying tasks. You can use `robocopy` to mirror entire directory trees, while preserving file attributes, file ownership, and NTFS permissions. Its extensive set of switches, such as `/mir` for mirroring directories and `/sec` for copying security settings, makes it an indispensable tool for system administrators and advanced users.

An excellent resource on how to use `robocopy` can be found at <https://learn.microsoft.com/en-us/windows-server/administration/windows-commands/robocopy>. The `robocopy` utility is the Swiss army knife of file and folder copy utilities. Let's look at some examples.

The `/mir` switch mirrors the ownership, attributes, and permissions from one directory to another. The following command-line command will mirror the contents of `C:\Source` to `D:\Destination`:

```
robocopy C:\Source D:\Destination /mir
```

The `/sec` switch ensures that files are copied along with their NTFS permissions, ensuring that the security settings are preserved. The following command-line command will copy all files from `C:\Source` to `D:\Destination` along with their NTFS security:

```
robocopy C:\Source D:\Destination /sec
```

The `/mov` switch moves files and deletes the old files from the source after copying. The following command-line command moves files from the source to the destination, deleting them from the source once the copy is complete:

```
robocopy C:\Source D:\Destination /mov
```

The `/xd` switch can exclude directories from being copied. The following command-line command will copy all files from the source directory to the destination directory. The command will also exclude the `$Recycle.bin` and `System Volume Information` directories.

```
robocopy C:\Source D:\Destination /  
xd $Recycle.Bin," System Volume Information"
```

The `/log:` switch allows you to specify a log file where `robocopy` will write its output, making it easier to review the results of the copy operation. The following command-line command will copy files from `C:\Source` to `D:\Destination` and log the output to `C:\Logs\robocopy.log`:

```
robocopy C:\Source D:\Destination /log:C:\Logs\robocopy.log
```

Network Connectivity Tools

Windows is a network operating system, which means that the operating system and its principal user rely on the network for connectivity to information. This is where the command line becomes really useful to the administrator of the PC. The command line will return a large amount of data that is normally not suited for a graphical user interface (GUI). The following is a short list of commands that can help you diagnose network connectivity issues from the command line.

***ipconfig* Command**

The `ipconfig` command is a network administrator's best friend—it assists in the diagnosis of network problems with the operating system. The `ipconfig` command without any switches displays basic information, such as the IP address, subnet mask, default gateway, and DNS suffix. The command `ipconfig /all` lists adapters and each one's assigned IP address, subnet mask, default gateway, DNS suffix, DNS server(s), DHCP server, and MAC address, just to name the most important elements. Viewing these assignments can help you diagnose the current network status of the connection.

In addition to verifying the status of a network connection's assignments, you can release and renew DHCP-assigned IP addresses. The `/release` switch releases the IP address, and the `/renew` switch renews the lease of an IP address.

The `ipconfig` command also allows you to view the local DNS cache with the `/displaydns` switch. You can flush the local DNS cache with the `/flushdns` switch. These switches come in handy when a DNS entry has changed, and you want to immediately flush the cache and verify any cached entries.

***ping* Command**

Next to the `ipconfig` command, the `ping` command is the runner-up as the network administrator's best friend. The `ping` command allows you to verify network connectivity via *Internet Control Message Protocol (ICMP)* packets. A common troubleshooting step used by network administrators is to ping the default gateway. If it returns a ping, then the network connectivity problem is probably beyond that device, or your subnet mask is incorrect. However, after a successful ping, you can verify that your computer has basic connectivity to it. An example of a successful ping is:

```
C:\Users\Sybex>ping 172.16.1.1
```

```
Pinging 172.16.1.1 with 32 bytes of data:
```

```
Reply from 172.16.1.1: bytes=32 time<1ms TTL=64
```

```
Reply from 172.16.1.1: bytes=32 time<1ms TTL=64
```

```
Reply from 172.16.1.1: bytes=32 time<1ms TTL=64
```

```
Reply from 172.16.1.1: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 172.16.1.1:
```

```
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

Approximate round trip times in milli-seconds:

Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Sybex>

tracert Command

The `tracert` command allows the network administrator to verify the path a network packet travels to its destination. As the diagnostic packet passes through the internetwork, each router responds with a response time. This enables you to locate a fault in an internet-work. Here's an example:

C:\Users\Sybex>tracert 8.8.8.8

```
Tracing route to google-public-dns-a.google.com      [8.8.8.8]
over a maximum of 30 hops:
  1    <1 ms    <1 ms    <1 ms  pfsense.Sybex.local  [172.16.1.1]
  2    13 ms    12 ms    17 ms  96.120.62.213
  3    15 ms    15 ms    12 ms  te04012.comcast.net  [68.86.101.141]
  4    13 ms    19 ms    12 ms  162.151.152.153
  5    12 ms    13 ms    20 ms  96.108.91.78
  6    22 ms    14 ms    20 ms  96.108.91.121
  7    21 ms    24 ms    20 ms  be-7016-cr02.comcast.net  [68.86.91.25]
  8    20 ms    20 ms    26 ms  be-10130-pe04.comcast.net  [68.86.82.214]
  9    20 ms    20 ms    21 ms  as040-2-c.comcast.net  [75.149.229.86]
 10    22 ms    21 ms    20 ms  108.170.240.97
 11    20 ms    23 ms    21 ms  108.170.226.85
 12    20 ms    22 ms    18 ms  google-public-dns-a.google.com  [8.8.8.8]
```

Trace complete.

C:\Users\Sybex>

pathping Command

`pathping`, another command-line tool, combines the benefits of `tracert` and `ping`. The tool can be used to diagnose packet loss (or suspected packet loss) to a destination website. It is invaluable to network administrators to help prove to their ISP that packet loss is a problem on their network.

The tool will first trace the entire path to a destination IP address or DNS host. Then, each of the hops will be tested with ICMP for packet loss and round-trip time. It easily identifies router hops that are causing the delay or packet loss. The following is an example of a `pathping` to a provider's DNS server:

C:\Users\Sybex>pathping 75.75.75.75

Tracing route to cdns01.comcast.net [75.75.75.75]

over a maximum of 30 hops:

```

0 Wiley.sybex.local [172.16.1.101]
1 pfSense.sybex.local [172.16.1.1]
2 96.120.62.213
3 te-0-5-0-12-sur02.pittsburgh.pa.pitt.comcast.net [69.139.166.77]
4 be-11-ar01.mckeesport.pa.pitt.comcast.net [68.86.147.109]
5 be-7016-cr02.ashburn.va.ibone.comcast.net [68.86.91.25]
6 ae-4-ar01.capitolhghts.md.bad.comcast.net [68.86.90.58]
7 ur13-d.manassascc.va.bad.comcast.net [68.85.61.242]
8 dns-sw01.manassascc.va.bad.comcast.net [69.139.214.162]
9 cdns01.comcast.net [75.75.75.75]

```

Computing statistics for 225 seconds...

```

Source to Here This Node/Link
Hop RTT  Lost/Sent = Pct Lost/Sent = Pct Address
0                               Wiley.sybex.local [172.16.1.101]
    0/ 100 = 0% |
1 0ms    0/ 100 = 0%  0/ 100 = 0% pfSense.sybex.local [172.16.1.1]
    0/ 100 = 0% |
2 14ms   0/ 100 = 0%  0/ 100 = 0% 96.120.62.213
    0/ 100 = 0% |
3 15ms   0/ 100 = 0%  0/ 100 = 0% te-0-5-0-12-sur02.pittsburgh.pa.pitt
.comcast.net [69.139.166.77]
    0/ 100 = 0% |
4 15ms   0/ 100 = 0%  0/ 100 = 0% be-11-ar01.mckeesport.pa.pitt.comcast.net
[68.86.147.109]
    0/ 100 = 0% |
5 23ms   0/ 100 = 0%  0/ 100 = 0% be-7016-cr02.ashburn.va.ibone.comcast.net
[68.86.91.25]
    0/ 100 = 0% |
6 23ms   0/ 100 = 0%  0/ 100 = 0% ae-4-ar01.capitolhghts.md.bad.comcast.net
[68.86.90.58]
    0/ 100 = 0% |
7 25ms   0/ 100 = 0%  0/ 100 = 0% ur13-d.manassascc.va.bad.comcast.net
[68.85.61.242]
    0/ 100 = 0% |
8 24ms   0/ 100 = 0%  0/ 100 = 0% dns-sw01.manassascc.va.bad.comcast.net
[69.139.214.162]
    0/ 100 = 0% |
9 23ms   0/ 100 = 0%  0/ 100 = 0% cdns01.comcast.net [75.75.75.75]

```

Trace complete.

C:\Users\Sybex>

netstat Command

The `netstat` command allows you to view listening and established network connections for the operating system. Several switches can be used with the `netstat` command. One of the most useful is the `-b` switch, which displays the name of the application and its current established connections. Adding the `-a` switch displays all the listening connections in addition to the established connections. A basic example follows:

```
C:\Users\Sybex>netstat
```

Active Connections

Proto	Local Address	Foreign Address	State
TCP	127.0.0.1:49750	view-localhost:50912	ESTABLISHED
TCP	127.0.0.1:50912	view-localhost:49751	ESTABLISHED
TCP	172.16.1.181:49208	104.20.60.241:https	ESTABLISHED
TCP	172.16.1.181:49599	172.67.181.149:https	ESTABLISHED
TCP	172.16.1.181:49600	52.167.17.97:https	TIME_WAIT
TCP	172.16.1.181:49602	20.50.80.210:https	ESTABLISHED
TCP	172.16.1.181:49603	a104-75-163-105:http	TIME_WAIT
TCP	172.16.1.181:56759	151.101.1.140:https	ESTABLISHED
TCP	172.16.1.181:64151	iad23s96-in-f10:https	CLOSE_WAIT
TCP	172.16.1.181:64152	iad66s01-in-f13:https	CLOSE_WAIT
TCP	172.16.1.181:64154	iad23s96-in-f10:https	CLOSE_WAIT

```
C:\Users\Sybex>
```

nslookup Command

DNS is one of the most important network services that an operating system and user relies on for resolution of `www.sybex.com` to an IP address. Without DNS, we just couldn't remember the millions of IP addresses; it would be like trying to remember the phone number of every person you've ever met or are going to meet.

When DNS problems arise, the `nslookup` command allows you to verify that DNS is working correctly and that the correct results are being returned. The simplest way to use DNS is to use an inline query, such as **`nslookup www.sybex.com`**. This will return the IP address associated with the *fully qualified domain name (FQDN)* of `www.sybex.com`. The `nslookup` command can also be used in the interactive mode by typing **`nslookup`** and pressing Enter. This mode allows you to query more than the associated IP address, depending on the type of DNS record you are trying to diagnose. By default, the record looked up with the `nslookup` command is the A or CNAME DNS records. These records are the most commonly looked-up DNS records for diagnosing connectivity issues. By specifying the `-type` argument, you can change the default record queried. The following is an example of retrieving the IP address for the FQDN of `www.sybex.com`, as well as the use of the `-type` argument:

```
C:\Users\Sybex>nslookup www.sybex.com
Server: pfsense.wiley.local
```


Address: 172.16.1.1

Non-authoritative answer:

Name: www.sybex.com

Address: 63.97.118.67

C:\Users\Sybex>nslookup -type=mx sybex.com

Server: pfSense.wiley.local

Address: 172.16.1.1

Non-authoritative answer:

sybex.com MX preference = 20, mail exchanger = cluster1a.us.message-labs.com

sybex.com MX preference = 10, mail exchanger = cluster1.us.message-labs.com

C:\Users\Sybex>

net Command

The **net** command can be used with several different subcommands. Most of the subcommands have been deprecated with the last few releases of Windows.

net use The **net use** subcommand is still widely used by administrators to map drive letters to network shares. The syntax for mapping a drive of Z: to the network location of \\server\share is as follows:

```
net use Z: \\server\share
```

net user **net user** is another subcommand widely used by administrators. This command allows an administrator to list all the local accounts on a Windows installation by entering **net user**. It can also be used to list local accounts or domain accounts by supplying some arguments to the command. If you wanted to create a local account of *usertwo* with a password of *Passw0rd*, you would use the following syntax:

```
net user usertwo Passw0rd /add
```

OS Management Tools

There are several other command-line tools just as useful as the connectivity tools we've discussed, but they are used for various other purposes. The following tools are part of the objectives for the 220-1202 exam, and they are commonly used by Windows administrators.

gpupdate Command

Active Directory refreshes local and Active Directory-based policies every 90 minutes in what is called a *background refresh cycle*. When the background refresh happens, policies are reapplied, forcing the settings that the administrator has configured in the Group Policy settings.

The `gpupdate` command is used to update Group Policy settings. It refreshes, or changes, both local and Active Directory–based policies and replaces some of the functionality that previously existed with the `secedit` command.

The `gpupdate` command can force the refresh cycle immediately with the `/force` switch. In addition, you can target the computer or the user, which is particularly useful when trying to diagnose a problem with Group Policy Objects (GPOs). The following is an example of forcing a refresh for the computer GPO settings:

```
C:\Users\bohack>gpupdate /force /target:computer
```

```
Updating policy...
```

Computer Policy update has completed successfully.

```
C:\Users\bohack>
```

***gpresult* Command**

The `gpresult` command is used to show the Resultant Set of Policy (RSOP) report/values for a remote user and computer. Bear in mind that configuration settings occur at any number of places; they are set for a computer, a user, a local workstation, the domain, and so on. Often one of the big unknowns is which set of configuration settings takes precedence and which is overridden. With `gpresult`, it is possible to ascertain which settings apply.

A number of switches can be used in conjunction with the `gpresult` command. The most useful switches are `/r` and `/z`. The `/r` switch allows you to see the RSOP summary of GPOs applied. This allows you to quickly verify if a policy is being applied. You can then use the `/z` switch to turn on super-verbosity, which allows the output to display the exact settings being applied.

***sfc* Command**

The System File Checker (`sfc.exe`) is a command line–based utility that checks and verifies the versions of system files on your computer. If system files are corrupted, `sfc` will replace the corrupted files with correct versions.

The syntax for the `sfc` command is as follows:

```
sfc [switch]
```

Table 3.6 lists the switches available for `sfc`.

TABLE 3.6 `sfc` switches

Switch	Purpose
<code>/scanfile</code>	Scans a file that you specify and fixes problems if they are found
<code>/scannow</code>	Immediately scans all protected system files and replaces corrupted files with cached copies

Switch	Purpose
<code>/verifyonly</code>	Scans protected system files and does not make any repairs or changes
<code>/verifyfile</code>	Identifies the integrity of the file specified and makes any repairs or changes
<code>/offbootdir</code>	Repairs an offline boot directory
<code>/offwindir</code>	Repairs an offline Windows directory

To run `sfc`, you must be logged in as an administrator or have administrative privileges. If the System File Checker discovers a corrupted system file, it will automatically overwrite the file by using a copy held in another directory. The most recent Windows versions store the files in a large number of discrete folders beneath `C:\WINDOWS\WINSXS` (where they are protected by the system and only TrustedInstaller is allowed direct access to them—the cache is not rebuildable). TrustedInstaller is a service in Windows 11 that enables the installation, removal, and modification of system components.



The `C:\WINDOWS\SYSTEM32` directory is where many of the Windows system files reside.

If you attempt to run `sfc` from a standard command prompt, you will be told that you must be an administrator running a console session in order to continue. Rather than opening a standard command prompt, on Windows 11 open the Start menu, type **cmd**, and then press **Ctrl+Shift+Enter**. The UAC will prompt you to continue, and then you can run `sfc` without a problem.

shutdown Command

The `shutdown.exe` utility can be used to schedule a shutdown (complete or a restart) locally or remotely. You can even have the computer enter into a hibernation power state. In addition, shutdowns of the computer can be logged with a variety of reasons, from unplanned to application maintenance. A message can also be specified and announced to users for the shutdown. The syntax of the command is as follows:

```
shutdown [/i | /l | /s | /sg | /r | /g | /a | /p | /h | /e | /o] [/hybrid]
[/soft] [/fw] [/m \\computer] [/t xxx] [/d [p|u:]xx:yy [/c "comment"]]
```

There are a lot of different switches, as you can see from the previous usage syntax. Table 3.7 lists the most important switches for the shutdown command.

TABLE 3.7 shutdown switches

Switch	Purpose
/s	Shut down the computer.
/sg	Shut down the computer. On the next boot, restart any registered applications.
/r	Do a full shutdown and restart the computer.
/g	Do a full shutdown and restart the computer. After the system is rebooted, restart any registered applications.
/a	Abort a system shutdown.
/h	Hibernate the local computer.
/o	Go to the advanced boot options menu and restart the computer. Must be used with the /r option.
/m \\computer	Specify the target computer.
/t xxx	Set the timeout period before shutdown to xxx seconds. The valid range is 0–315360000 (10 years), with a default of 30. If the timeout period is greater than 0, the /f parameter is implied.

Disk Commands

There are several commands for the administration of drives and folders. In this section, we will cover the common command-line tools that you will use in day-to-day administration of Windows. All of these tools can be accessed in the GUI as well, but the advantage to using the command line is avoiding clicks and potential mistakes.

format Command

The `format` command is used to wipe data off disks and prepare them for new use. Before a hard disk can be formatted, it must have partitions created on it. (Partitioning was done in the DOS days with the `fdisk` command, but that command does not exist in current versions of Windows, having been replaced by `diskpart`.) The syntax for `format` is as follows:

```
format [volume] [switches]
```

The *volume* parameter describes the drive letter (e.g., D:), mount point, or volume name. Table 3.8 lists some common `format` switches.

TABLE 3.8 format switches

Switch	Purpose
<code>/fs:[filesystem]</code>	Specifies the type of filesystem to use (FAT, FAT32, or NTFS)
<code>/v:[label]</code>	Specifies the new volume label
<code>/q</code>	Executes a quick format

There are other options as well—to specify allocation sizes, the number of sectors per track, and the number of tracks per disk size. However, we don't recommend that you use these unless you have a very specific need. The defaults are just fine.

Thus, if you wanted to format your D: drive as NTFS, with a name of HDD2, you would type the following:

```
<?b Start?>format D: /fs:ntfs /v:HDD2<?b End?>
```



Before you format any drive, be sure that you have backed it up (or are prepared to lose whatever is on it).

diskpart Command

The `diskpart.exe` utility shows the partitions and lets you manage them on the computer's hard drives. You can perform the same functions in the `diskpart` utility as you can perform in the GUI, which is discussed later in this chapter. Because of the enormous power that `diskpart` holds, membership in the Administrators local group (or equivalent) is required to run `diskpart`.

chkdsk Command

You can use the Windows `chkdsk.exe` utility to create and display status reports for the hard disk. `chkdsk` can also correct filesystem problems (such as cross-linked files) and scan for and attempt to repair disk errors. You can manually start `chkdsk` by right-clicking the problem disk and selecting Properties. This will bring up the Properties dialog box for that disk, which shows the current status of the selected disk drive.

By clicking the Tools tab at the top of the dialog box and then clicking the Check button in the Error-checking section, you can start `chkdsk`. Exercise 3.2 walks you through starting `chkdsk` in the GUI, and Exercise 3.3 does the same from the command line.

EXERCISE 3.2**Running chkdsk within Windows**

1. Open Windows File Explorer by holding down the Windows key and pressing E.
2. Right-click C: and choose Properties.
3. Click the Tools tab, and then click the Check button.
4. Choose Scan Drive.
5. Once the drive is scanned for errors, you can click Close.

In operating systems prior to Windows 10, you can choose options such as Automatically Fix Filesystem Errors and/or Scan For And Attempt Recovery Of Bad Sectors.

EXERCISE 3.3**Running chkdsk at the Command Line**

1. Open an administrative command prompt in Windows 10/11 by clicking the Start menu and typing **cmd**, and then press Ctrl+Shift+Enter.

You can also right-click on the result and run **cmd** with administrator privileges from the context menu.

2. Type **chkdsk /f** and press Enter.

The system will now scan for, and fix, filesystem errors.

Informational Tools

Command-line tools like **winver** and **hostname** are essential utilities for quickly accessing system information in Windows environments. These tools, while simple, are invaluable for IT professionals and users who need to efficiently gather system details without navigating through multiple settings menus. The following are common informational tools you need to know for the 220-1202 exam.

hostname Command

The **hostname** command allows the administrator to keep their sanity. The command returns the hostname of the computer that you have the command prompt open on. It can get pretty confusing for the administrator when they jump from one computer to another and remain in the command line. So by typing the command **hostname**, you can positively identify the system you are about to execute a command on. The following is an example of the command's use:

```
C:\Users\Sybex>hostname
```

Wiley-023432

C:\Users\Sybex>

winver Command

`winver.exe` is not a command-line tool but a useful tool to glean information about the operating system. The command will display a GUI dialog box. The command's sibling at the command line is `ver.exe`. This command has its roots all the way back to DOS, and it's still active in the latest version of Windows 11. Using the command will return the version of Windows, as shown in the following:

C:\Users\NetworkedMinds>`ver.exe`

Microsoft Windows [Version 10.0.19043.1237]

C:\Users\NetworkedMinds>

whoami Command

The `whoami` command in Windows is a simple yet powerful tool used to display the current user's username. When executed in the command prompt, it returns the domain and username of the account currently logged into the system. The following are some of the possible switches that can be used with the `whoami` command.

To simply find out the username of the account currently logged in, you can use the `whoami` command. This is particularly useful when you need to confirm which user is active on the system, especially in environments with multiple user accounts. The following command will display the current user's domain and username:

C:\Users\NetworkedMinds>`whoami`

DOMAIN\username

If you need to check what privileges the current user has, the `whoami /priv` command is very handy. This can help you understand what actions the user is authorized to perform, which is crucial for troubleshooting permission issues. The following command lists the privileges assigned to the current user, which can be useful for understanding what actions the user can perform:

C:\Users\NetworkedMinds>`whoami /priv`

PRIVILEGE_NAME	DESCRIPTION	STATE
SeShutdownPrivilege	Shut down the system	Disabled
SeChangeNotifyPrivilege	Bypass traverse checking	Enabled
SeUndockPrivilege	Remove computer from docking station	Disabled

To see all the groups the current user belongs to, you can use the `whoami /groups` command. This is useful for verifying the user's roles and permissions within the system, as group memberships often determine access rights. The following command shows all the groups the current user is a member of, providing insight into the user's permissions and roles:

```
C:\Users\NetworkedMinds>whoami /groups
```

GROUP_INFORMATION

```
=====
```

*Group Name	Attributes
DOMAIN\Domain Users	Mandatory group, Enabled by default, Enabled group
BUILTIN\Administrators	Mandatory group, Enabled by default, Enabled group

When you need to find the Security Identifier (SID) for the current user, the `whoami /user` command is the way to go. The SID is a unique value used to identify the user account, which can be important for various administrative tasks. The following command returns the Security Identifier (SID) for the current user, which is a unique value used to identify the user account:

```
whoami /user
```

USER_INFORMATION

```
=====
```

User Name	SID
DOMAIN\username	S-1-5-21-3623811015-3361044348-30300820-1013

help and /? Commands

The `help` command does what it says: It gives you help. If you just type **help** and press Enter, your computer gives you a list of system commands that you can type. To get more information, type the name of a command that you want to learn about after typing **help**. For example, if you type **help rd** and press Enter, you will get information about the `rd` command.

You can get the same help information by typing `/?` after the command.




The `/?` switch is slightly faster and provides more information than the `help` command. The `help` command provides information for system commands only. (It does not include network commands.) For example, if you type **help ipconfig** at a command prompt, you get no useful information (except to try `/?`); however, typing **ipconfig /?** provides the help file for the `ipconfig` command.

Commands with Standard Privileges vs. Administrative Privileges

By default, any user can open a command prompt and begin typing the names of command-line commands. Certain commands, however, can be dangerous when run and, as a safety precaution, require administrative privileges. The `sfc` command was mentioned earlier, for example, as requiring administrative privileges.

With Windows 11, rather than opening a standard command prompt, choose Start menu > All > Windows Tools, right-click Command Prompt, and then choose Run As Administrator. The UAC will prompt you to continue, and then you can run `sfc` without a problem, as shown in Figure 3.39.

FIGURE 3.39 Opening a command prompt with admin privileges in Windows 11



In most cases, if you try to run a utility that requires administrative privileges and you are not currently in a console session that has the privileges, an error message will notify you of this.

Networking in Windows

CompTIA expects you to know a number of topics related to networking and Windows. This section covers the various scenarios in which you will deploy Windows in a network. First, we'll cover small office, home office (SOHO) deployments, and then we'll scale it out to an enterprise network.



Exam Objective 1.7, “Given a scenario, configure Microsoft Windows networking features on a client/desktop,” requires you to know various network components of Windows 11. Be sure to know the various network models and their differences, the various ways to connect to a network, access resources, and configure a client.

Networking Models

You can choose from several different networking models to facilitate Windows authentication. There is no single way to implement a network model; it all depends on the needs of the organization to share the resources. Each model has its advantages and disadvantages. This section covers the three most popular methods.

Workgroups

The *workgroup* networking model has existed in the Windows operating system since it was first introduced. Windows can function as both a client and a server simultaneously. When an operating system can function as both a client and a server, it is considered *peer-to-peer networking*. Clients can join the network and leave the network at any time.

The core of the Windows operating system is essentially the same for both a server and workstation. There are a few differences between the two—for instance, servers normally give priority to background services like file and printer sharing versus the desktop. The second and most limiting is the restriction of connecting 20 simultaneous clients to a single workstation. After 20 simultaneous connections are made, all others clients will return an Access Denied error—that is, until one of the 20 connections is closed out and the session is terminated. This takes some time on the workstation side, so if you have more than 20 clients, you should plan to deploy an actual server system running software such as Windows Server 2022 or Windows Server 2025.

Workgroups are normally used in SOHO environments or in situations that do not require the infrastructure of a dedicated server for authentication. They should be kept to a maximum of 20 clients, with the expectation that each client will maintain its own resources (files and printer sharing). Many small offices use this networking model and never need anything more.

A typical situation where a workgroup is effective is when a printer needs to be shared from a single computer. The disadvantage is that the computer must be on in order for the

clients to use the printer. Another disadvantage with workgroups is user authentication, which we will discuss in detail later in this chapter.

When you install Windows 11, by default it is joined to the workgroup named Workgroup. However, there may be instances in which you want to join another workgroup. To join another workgroup, perform the following steps:

1. Click the Start menu.
2. Select Settings.
3. Click System.
4. Click About.
5. Click the Advanced System Settings, Related Setting.
6. Select Change Settings under the heading of Computer Name, Domain, And Workgroup Settings.
7. Click Change Next To Rename This Computer or Change Its Domain Or Workgroup.
8. Change the Workgroup option to the desired workgroup.
9. Click OK.
10. Click OK.
11. Reboot the computer.


Domains

Domain functionality has existed since Microsoft Windows 3.51 (mid-1990s). Unlike the workgroup networking model, the domain networking model requires that clients be joined to a domain. Joining a domain creates a trust between the client (Windows client, for example) and the authentication server (Windows Server 2025 running Active Directory, for example). Joining a client to a domain allows users with an account in the domain to log into the client. To join a domain in Windows 11, perform the following steps:

1. From the Start menu, choose Settings > Accounts > Access Work Or School > Connect > Join This Device To A Local Active Directory Domain. You will be prompted to enter a domain name, as shown in Figure 3.40.
2. Enter your credentials on the domain when prompted, which allows the joining of the operating system, as shown in Figure 3.41.
3. Reboot the operating system for the changes to take effect.

An alternate way of joining a domain is through the System Properties dialog box. To do so:


1. Right-click the Start menu.
2. Click System.
3. Under Related Links choose Advanced System Settings.
4. Under Computer Name, Domain, And Workgroup Settings, select Change Settings.

FIGURE 3.40 Windows 11 Settings app for joining a domain

5. Click Change.
6. Select the Domain radio button under Member Of.
7. Enter the domain name and click OK, as shown in Figure 3.42.

The next dialog box will prompt you to enter your credentials on the domain that allows the joining of the operating system. You will then need to reboot the operating system for the changes to take effect. This method of joining an operating system to a domain has been supported since Windows XP.


Domains also allow for files and printers to be secured with domain credentials. The key takeaway in the benefit to domains is centralized authentication for users and computers. It is important to note that the term *domain* describes both the networking model and the friendly name of the security domain of the network.

FIGURE 3.41 Windows 11 domain credentials prompt

Microsoft Active Directory is the main authentication technology used with domain controllers in a domain networking model. Active Directory contains objects such as users, computers, and printers, as well as many other types of objects, called Group Policy Objects (GPOs). Active Directory allows for the grouping of these objects logically so that they can be controlled with policies through Group Policy. Active Directory domains can scale to any number of joined clients and user credentials. Many large organizations use Active Directory for centralized user authentication, because it is very scalable.



Group Policy is a component of Active Directory as well as a component of the client-side operating system. Group Policy is used to create policies that can be applied to groups of users and a group of computers. Group Policy is considered enforced, since it reapplies policies every 90 minutes.


FIGURE 3.42 Windows domain joining

Microsoft Entra ID

Microsoft Entra ID, previously known as Azure Active Directory, is a cloud-based service that helps manage user identities and access to resources. It allows users to sign in once and access multiple applications securely. Microsoft Entra ID has features like single sign-on (SSO), multifactor authentication (MFA), and conditional access policies. These features ensure only authorized users can access important data. Entra ID also integrates easily with various applications and services, making it simpler for IT administrators to manage user permissions and keep everything secure. With the explosion of cloud-based applications, just about every application supports Microsoft Entra ID. Microsoft 365 exclusively uses Microsoft Entra ID to secure information and provide identity services to users.

To join a Windows 11 device to Microsoft Entra ID, follow these steps:

1. Click Start menu > Settings > Accounts > Access Work Or School > Connect > Join This Device To Microsoft Entra ID.
2. Enter your Microsoft Entra ID email address and click Next, as shown in Figure 3.43.
3. On the next screen enter your password, and click Sign In.

FIGURE 3.43 Entra ID Sign In

4. You will need to enter your authenticator code or satisfy MFA on the next screen, then click Verify.
5. On the next screen you will verify the organization and the username that you will join with, then click Join, as shown in Figure 3.44.
6. The last screen welcomes you to the Microsoft Entra ID organization, and you can click Done.

Once authenticated, the device will be registered with Microsoft Entra ID, and any applicable policies will be applied automatically. You can verify the organization and user you are connected as, on the Access Work Or School screen, as shown in Figure 3.45. The device is now enrolled in the organization's mobile device management (MDM) solution of Intune. This enrollment allows administrators to manage the device remotely, deploy applications, and enforce security settings. The integration with MDM solutions ensures that devices remain compliant with organizational policies and can be monitored for any security issues. This is all dependent on the subscription purchased.



Microsoft Entra ID is not a CompTIA objective, but it should be. More and more organizations are adopting Microsoft Entra ID. It was included in this book for completeness. You can find more information by going to www.microsoft.com/en-us/security/business/identity-access/microsoft-entra-id.

FIGURE 3.44 Make Sure This Is Your Organization






FIGURE 3.45 Access Work Or School



User Authentication

When you choose to use a workgroup networking model, all authentication is local to the operating system. Windows operating systems contain a local authentication mechanism called the *Security Account Manager (SAM)*. The SAM can be considered a local database of users and groups. All users locally authenticating to the workstation authenticate against this internal database of usernames and passwords and are granted a local access token, as shown in Figure 3.46. The local access token allows the user to access local resources secured with the user's identity.

FIGURE 3.46 Windows local authentication




When clients are joined to a domain, a user with credentials on the domain can log into the workstation. When this happens, they are authenticating against an Active Directory *domain controller*. An Active Directory domain controller retains information about all access rights for all users and groups in the network. When a user logs into the system, AD issues the user a *globally unique identifier (GUID)*, also known as an *access token*. Applications that support Active Directory can use this access token to provide access control. When a client is joined to a domain, users can still log into the local operating system using the SAM, as if they have a local account (see Figure 3.47). However, local logins are normally restricted to administrators once a client is joined to the domain.

Active Directory simplifies the sign-on process for users and lowers the support requirements for administrators. Access can be established through groups, and it can be enforced through group memberships. Active Directory can be implemented using a Windows Server

(such as Windows Server 2022 or Windows Server 2025) computer. All users will then log into the Windows domain using their centrally created AD accounts.

FIGURE 3.47 Windows domain authentication



Microsoft Entra ID is a cloud-based alternative to hosting a traditional AD domain in a physical location. There are different benefits, such as the Microsoft 365 platform, depending on the requirements. There is also a monthly cost to maintain the licensing.

One of the big problems that larger systems must deal with is the need for users to access multiple systems or applications. This may require a user to remember multiple accounts and passwords. The purpose of a *single sign-on* (SSO) is to give users access to all the applications and systems that they need when they log in.

SSO operates on the principle that the resource trusts the authentication server. When a user logs in initially, they will authenticate against the authentication server for their organization. When the user then visits the resource, which is normally a cloud-based resource, it will prompt the authentication server to provide a claim on behalf of the user, as shown in Figure 3.48.


The claim normally contains basic information about the user, such as first and last name, email address, or any other attribute. At no time is the user's password sent, because they authenticated once already. Although we've oversimplified SSO in this example, it really is this simple, without the layers of encryption and complicated trust rules. As we adopt more and more cloud resources, it is becoming the number one way to provide authentication for our users because we never transmit the actual username and password.

Establishing a Network Connection

A key element to a successful network is the connection that connects the computer to the actual network. There are a number of different ways to connect to the network that we will

cover in the following section. The key takeaway is that the network system will function identically, regardless of the connection. For example, if you have a computer joined to an organization's domain, the login process will function identically with a wireless versus a wired connection.

FIGURE 3.48 Single sign-on



The type of connection you choose is based on your convenience and the requirements of the connection. As an example, wireless is extremely convenient, but it requires you to be within distance of an access point (AP), and the further away you are from the AP, the slower the connection. If your requirement is consistent high bandwidth for a voiceover IP (VoIP) call, then a wired connection is your best choice and wireless would be problematic. The following are the most common network connection methods:

Wired A wired network is most common in organizations with desktop computers. Unfortunately, a wired connection means that there is no mobility for the connection; you are literally wired to a desk.


Wired connections are the most reliable and arguably easiest to diagnose when there are problems. The network link light gives you a visual indicator that you have a network connection. Windows 11 will also place a visual notification of a computer with a cable in the notification area when a wired connection is detected.

Wireless A wireless connection is found in networks where mobility is required. Wireless is often used in small office/home office (SOHO) network settings. It can also be found in small and large organizations that require workers to move around, such as a factory setting, sales workforce, or medical setting, just to name a few.

Connecting to wireless networks is not as straightforward as making a wired connection. There is rarely a visual identifier that you have a wireless connection, such as the link light

with a wired connection. Windows 11 will display the notification tray icon as a radio wave. If the radio wave is grayed out with an asterisk at the upper left, then there are wireless connections detected and you are not connected to any of them. If you click the wireless icon, all of the wireless networks available will be displayed, along with their security status, as shown in Figure 3.49.

FIGURE 3.49 Wireless connectivity



You can then choose a wireless network and select **Connect**. If the wireless network requires additional security, such as a preshared key (PSK) or a corporate login to a *captive portal*, the operating system will direct you. By default, wireless networks that you connect to will automatically reconnect when you are in range. When the wireless connection is established, the notification tray icon will appear as a white radio wave. As you move further away from the AP, the wireless indicator will act as a signal strength meter.


Virtual Private Network (VPN) In recent years, *virtual private networks (VPNs)* have become the rage of network privacy for browsing. These applications provide encryption to an anonymous server from which you can browse the Internet. However, the 220-1202 exam is based on traditional VPN technology that provides a secure connection between the two endpoints in an organization.

The VPN connection is an overlay network on top of an established network connection called the underlay network. So, you will need an established Internet connection before a VPN can connect. The VPN connection will provide two distinct features to the end user. The first feature it provides is an entry point to the organization's private network.

The second feature it provides is end-to-end encryption for everything transmitted over the connection between the client and the organization's private network.

Establishing a VPN connection will require information from the organization's VPN appliance or server. You will need the VPN protocol, server address, and the sign-in info. You can create a VPN connection by choosing the Start menu > Settings > Network & Internet > VPN; from here you can add a VPN connection, as shown in Figure 3.50.

FIGURE 3.50 VPN connectivity




Wireless Wide Area Network (WWAN) A wireless wide area network (WWAN) connection is a connection that is created with a cellular data provider, such as Sprint, Verizon, or AT&T, just to name a few. The mobile device will require a special card called a WWAN adapter, or most mobile devices will have a built-in card. You will then need to register the connection with the cellular provider to activate the connection; this process usually requires a monthly billing for data, and the connection is often metered.

A *metered connection* is a connection in which you pay for a specific amount of data to be downloaded. Once the amount of data is reached, you generally pay for overages per gigabyte. Many features in Windows, such as Windows Updates, will work differently

depending on whether the connection is a metered connection. Windows Updates will not download over a metered connection, so the features don't use all your precious data over the metered connection. The cellular option of connectivity is not the only connection that can be classified as a metered connection; each of the methods of connectivity in this section can be set as a metered connection. It is just more common to find a cellular connection as a metered connection.

Once the connection is registered, you can connect via the wireless notification tray icon and select the cellular network you wish to connect to. You can also connect to the cellular network by choosing the Start menu ► Settings ► Network & Internet ► Cellular, as shown in Figure 3.51. If your device does not support a WWAN device, then the Cellular section will not appear on the Network & Internet screen. After selecting Cellular, you can then configure roaming options, as well as the preference of cellular over Wi-Fi, and you can tell Windows to treat the connection as a metered connection.

FIGURE 3.51 Windows Cellular




Proxy Settings In addition to establishing a connection, you may need to set a proxy server depending on your organization's policies. In many organizations, the web browser is not allowed to directly request web pages from the destination web server. An intermediary called a *proxy server* is used to request the web page on behalf of the user.

The use of a proxy server allows for caching of frequently accessed web pages, as well as the ability to filter content. The proxy is primarily for web-based traffic, such as browsing with the Edge browser or Internet Explorer. However, other applications can also elect to use the proxy server, depending on their traffic type.

To configure the proxy settings for Microsoft Edge and Internet Explorer, from the Start menu, choose Settings > Network & Internet > Proxy. From the proxy screen you can configure the operating system to automatically use a setup script (JavaScript) by clicking the switch for Use Setup Script and specifying the script address, as shown in Figure 3.52.

FIGURE 3.52 Windows Proxy settings



You can also specify a manual proxy setup, which is a common configuration task. You will simply click the switch Use A Proxy Server and then enter the address and port. Specific websites often require direct access and will not work with a proxy server. You can enter exceptions into the lower section and separate servers with a semicolon. You can also use wildcards if you want to exclude an entire namespace.

Accessing Resources


Now that you understand networking models, authentication, and how to connect to the network, let's focus on the resources that users will access. This section covers the most common types of resources.

Network Shares

A *network share* is a type of resource sharing that allows access to file and folder resources over a network from a file server. File Explorer is used to access network shares, but you can also use command-line commands to access the resources, as you will see in the following sections. As the name implies, network shares are those that exist on the network; however, they can be mapped to appear as if they are local. The `net use` command can be used to establish network connections at a command prompt, for example. If you want to connect to a shared network drive and make it your M: drive, the syntax is **`net use M: \\server\share`**. The **`\\server\share`** portion of the command is called the *Universal Naming Convention (UNC)* path. The UNC path is a standard way to describe the server and file-share to the Windows operating system.

In addition to using the command line, you can use the GUI to map a network share. After browsing to the server by typing **`\\servername`** in the Windows File Explorer address bar, right-click the fileshare and select Map Network Drive. If you are using Windows 11, you will need to click Show More Options first. You will then be prompted with some options for mapping the network share, as shown in Figure 3.53.

FIGURE 3.53 Mapping a network drive





By default, the Reconnect At Sign-In option is selected. In order to ensure that the mapped drive is reconnected when the user signs in, this option must be selected.


Printer Shares

Another common type of network resource is printing. Although we often need to access files and folders, at some point we will probably need to print. You can use the same command to connect to a shared printer; the syntax of `net use lpt1: \\server\printername` will map a printer to the LPT1 device. As in the prior example, you can make the printer act as if it is locally connected to the operating system. You can also use the GUI method of connecting a printer by right-clicking the printer after browsing to the server and then selecting Connect, as shown in Figure 3.54.



Sharing printers was one of the primary reasons why networking became popular in the workplace, and it still continues to be an important reason today. Know that you can map to any printer on the network, but it still must be shared in order for you to be able to use it.

FIGURE 3.54 Connecting to printers



Administrative Shares

Administrative shares are automatically created for administrative purposes on Windows. These shares can differ slightly based on which operating system is running, but they end with a dollar sign (\$) to make them hidden. There is one for each volume on a hard drive (C\$, D\$, and so forth), as well as admin\$ (the root folder—usually C:\WINDOWS), and print\$ (where the print drivers are located). These are created for use by administrators and usually require administrator privileges to access. It's important to note that they are hidden shares and that any shared folder created with a trailing \$ will be hidden from the users as well. Unless you know they exist, they will not be visible.

Shared Resources

Although you can access shares on file servers and printers that are shared, there are many other shared resources that can be shared and accessed. Files and printers are just the common resources that make up the file and print sharing service that Windows has had built in, arguably since the inception of MS-DOS. Application sharing is one example of a shared resource. Scanners and faxes are other examples of hardware-based shared resources. These examples only scratch the surface; the cloud is full of applications and shared resources. All these resources can be authenticated with domain-based authentication or separate locally based credentials inside the application.

Firewall Settings

Microsoft Windows has come with a preinstalled firewall since Windows XP Service Pack 2. The addition of the Windows Firewall feature was welcomed at a time in history when the Internet could be described as the *Wild West*. Today Windows Defender Firewall is an integral part of Windows. Surprisingly the interface has not changed all that much from its debut with Windows XP. The original firewall was not turned on by default, but with the introduction of Windows Vista the firewall was on by default in the inbound direction. The firewall is never configured in an outbound direction, but if needed it has the ability.

Windows Defender Firewall is scalable as a host-based firewall, because an average user can configure the firewall. However, if more complicated firewall rules need to be composed, the advanced interface allows for an administrator to intervene. The configuration of the firewall is really straightforward; you can access the basic firewall controls by choosing from the Start menu Settings > Privacy & Security > Open Windows Security > Firewall & Network Protection. This will open the basic Windows Defender Firewall dialog, as shown in Figure 3.55. This dialog box allows you to perform basic firewall tasks, such as turning off the firewall, changing user notification, restoring defaults, and most importantly, allowing an app or feature through the firewall.

A number of applications and services are preconfigured in the firewall. For example, when you share a folder, the ports associated with filesharing are automatically enabled. Another mechanism exists to allow the firewall to easily configure itself; when a program is launched that listens to a port, an Allow or Cancel notification is sent to the user. If you click Show More,

you can select which firewall profile to apply the firewall exception, as shown in Figure 3.56. If the user selects Allow, a rule is added to the firewall for the specific application.

FIGURE 3.55 Windows Defender Firewall






FIGURE 3.56 Windows Defender Firewall notifications



Along with the automated mechanism in which firewall rules can be added, you can specify which network profile they are active in. Network profiles are identified by the MAC address of the default gateway. The firewall will learn the internal network of your home by making a note of the MAC address of your router. If the MAC address has not been seen before by the network operating system, then the firewall will ask if the network is private or public. This way the firewall can behave differently in your home than if you were in an airport or another public setting. There are three different network profiles that firewall rules can be active for: public, private, and domain. You can control if your laptop is public or private, but if a domain controller exists and the laptop is joined to a domain, then the domain profile becomes active.

You can get a more granular view of the firewall by launching the Windows Defender Firewall with Advanced Security MMC. To launch the MMC, from the Start menu choose All ► Windows Tools ► Windows Defender Firewall With Advanced Security. The individual rules can be examined or individually configured with the MMC, as shown in Figure 3.57. You can see each of the rules along with its effective profile of public, private, domain, or all.

FIGURE 3.57 Windows Defender Firewall with Advanced Security



Adding rules manually allows for maximum granularity but comes with the price of complexity. You can add a rule based on a program, port, predefined rule, or something totally custom. For example, if you wanted to only allow an incoming port of 2233 via TCP to a specific application awaiting its request for a specific network profile, this interface will allow you to do so. If you are the administrator of a domain, you can also create rules inside a Group Policy Object (GPO) and deploy the rules out to a large group of computers.



A Group Policy Object (GPO) is a collection of Group Policy settings. GPOs allow an administrator to change the behavior and appearance of the operating system and control certain aspects of the user.

Interface Firewall Profile

Windows Defender Firewall uses network profiles to apply different security settings based on the type of network you're connected to. This ensures that your device is appropriately protected whether you're at home, at work, or on the go. The mechanism for detecting the network firewall profile to use is based off the gateway's MAC address and your prior answers of private or public. Keep in mind that a private firewall profile will allow certain resources, whereas a public firewall profile will continue to block access to certain resources. The following are the three type of firewall profiles; you can control only two of them: private and public.

Private Network Profile The Private network profile is tailored for environments where you trust the other devices on the network. This could be your home, where you have control over the network and the devices connected to it, or your workplace, where security measures are typically in place. When your device is set to a private network profile, it allows for network discovery, meaning your device can find and interact with other devices on the same network. This is useful for sharing files, using network printers, and other collaborative tasks. Essentially, the private profile assumes a higher level of trust and facilitates easier communication between devices.

Public Network Profile On the other hand, the Public network profile is designed for use in less secure environments, such as coffee shops, airports, or any public Wi-Fi network. In these settings, the risk of encountering malicious devices or individuals is higher. Therefore, the public profile disables network discovery and file sharing features to keep your device hidden from others on the same network. This minimizes the risk of unauthorized access and helps protect your personal information and data from potential threats.

Domain Network Profile The Domain network profile is used when your device is connected to a network with a domain controller, typically found in corporate environments. This profile is automatically applied when your device can authenticate to a domain controller. The domain profile allows for more relaxed security settings compared to the public profile, as it assumes that the network is managed and secured by IT professionals.

By understanding the differences between these profiles and knowing how to switch between them, you can ensure that your device is always using the appropriate security settings for your current environment. To change the profile for the active connection follow these steps:

1. From the Start menu, choose Settings > Network & Internet.
2. Select the interface of either Ethernet or Wireless.
3. Set Network Profile Type with the Public Network and Private Network radio boxes.

Client Network Configuration

If you have a router or a server that supports the Dynamic Host Configuration Protocol (DHCP), the client will automatically configure itself with an IP address, subnet mask, default gateway, and the appropriate Domain Name System (DNS) servers. This is the default behavior for all devices, because it is trouble-free from a user's point of view. For example, if you turn on your laptop at home, your router will serve all the information necessary to get on the home network and the Internet. If you pick up your device and go to work, this process will also happen at your workplace to allow you to connect to servers or the Internet.

Elements of Client Configuration

Configuring a Windows client involves several critical elements that ensure network connectivity and efficient communication between devices. Understanding these elements not only helps in troubleshooting network issues but also enhances overall network performance and security. Let's look at each of the key elements in a Windows client configuration:

Internet Protocol (IP) Addressing Scheme An Internet Protocol (IP) addressing scheme is a systematic method used to assign unique addresses to devices on a network. In a Windows client configuration, IP addresses can be assigned either statically or dynamically. Static IP addressing involves manually setting a fixed IP address for a device, ensuring it remains constant over time. Dynamic IP addressing, on the other hand, uses DHCP to automatically assign IP addresses from a pool, allowing for efficient management and reuse of addresses.



IP addressing was covered in detail in Volume 1 Chapter 7, "Introduction to TCP/IP."

Domain Name System (DNS) Settings Domain Name System (DNS) settings in a Windows client configuration are essential for translating human-readable domain names into IP addresses that computers use to identify each other on the network. Proper DNS configuration ensures that when a user types a web address into their browser, the DNS server quickly resolves it to the corresponding IP address. Windows clients can be configured to use specific DNS servers, either provided by the Internet service provider (ISP) or custom servers like Google DNS or OpenDNS, to enhance speed, reliability, and security of domain name resolution.

Subnet Mask A subnet mask is a critical component in a Windows client configuration that defines the network and host portions of an IP address. It works in conjunction with the IP address to determine which part of the address identifies the network and which part identifies the individual device. For example, a common subnet mask of 255.255.255.0 indicates that the first three octets represent the network, while the last octet represents the host.

Gateway The gateway in a Windows client configuration acts as an intermediary device that routes traffic from the local network to other networks, including the Internet. It serves as the exit point for data leaving the local network and the entry point for data coming into the local network. Configuring the correct gateway address ensures


that devices on the network can communicate with external networks and access online resources. Typically, the gateway is the IP address of a router or firewall that connects the local network to the broader Internet.

Configuration of a Static IP Address

If you need to statically configure your network settings, that will require some planning and manual configuration. The first thing you will need is an IP address that is not used by another computer in the network. The subnet mask will also need to match the network you will configure the computer in. If you want to communicate outside the immediate network, you will need a default gateway, which is your router's IP address. A DNS address is also required if you want to translate simple domain names to IP addresses. There are two ways to configure the static IP address: by using the new Settings app and by using the legacy Control Panel applet. You should be familiar with both ways, as the legacy Control Panel applet offers more features, such as alternate IP address configuration.

Settings App You can configure the network addressing via the Settings app by clicking Start menu > Settings > Network & Internet > Properties, and then clicking Edit under IP Assignment. You can then choose Manual from the drop-down menu to open two switches for IPv4 and IPv6. You can then click the switch for IPv4 and enter the IP address, subnet mask, gateway, and DNS settings, as shown in Figure 3.58.

FIGURE 3.58 Settings App network configuration




Control Panel Applet The original Control Panel applet has been around since Windows NT 4.0, and you can get there a number of ways. To open the dialog, first open Control Panel, then click Network And Sharing Center. When the dialog box opens, select Change Adapter Settings from the left side and then right-click the network adapter and select Properties. You can then double-click Internet Protocol Version 4 (TCP/IPv4). From here you can select Use The Following IP Address and manually enter the IP address, subnet mask, and gateway, as shown in Figure 3.59. As you can see, this method is not easier to get to, but it is familiar to a seasoned professional.

Configuring an Alternative IP Address in Windows


Windows 11 and prior versions of Windows also allow for the use of an alternate IP address—that is, an address configured for the system to use in the event the first choice is not available. In order for an alternate configuration to be set, the first choice has to be dynamic; the tab becomes visible only when the General configuration is set to Obtain An IP Address Automatically (as shown in Figure 3.59), and the alternate is used only if the primary address cannot be found/used, such as when the DHCP server is down. The Alternate IP address configuration is only available with the Control Panel applet. The following is a list of the elements you can configure for an alternative IP address:

FIGURE 3.59 Control Panel network configuration



IP Addressing Two radio buttons exist on the Alternate Configuration tab: Automatic Private IP Address and User Configured (see Figure 3.60). The first option is the default, meaning that the alternate address used is one in the APIPA range (169.254.x.x). Selecting User Configured requires you to enter a static IP address to be used in the IP Address field. The entry entered must be valid for your network in order for it to be usable.

FIGURE 3.60 Alternate Configuration tab



Subnet Mask When the User Configured radio button is selected on the Alternate Configuration tab, you must enter a value in the Subnet Mask field. This value must correspond with the subnet values in use on your network and work with the IP address you enter in the IP Address field.

DNS When the User Configured radio button is selected on the Alternate Configuration tab, you should enter values in the Preferred DNS Server and Alternate DNS Server fields. These entries are needed in order to translate domain names into IP addresses.

Gateway When the User Configured radio button is selected on the Alternate Configuration tab, you must enter a value in the Default Gateway field. This value must correspond with the subnet values and the IP address that you enter. This address identifies the router to be used to communicate outside the local network.

Summary

In this chapter, you learned the various options for the installation and upgrade of Windows. Both the installation and the upgrade process were covered in great detail, so you could see what happens in each step. We also covered the various ways to install Windows, deploy images, and recover Windows when things go wrong.

In addition, you learned some of the command-line tools that can be used to administer Windows. We covered basic Windows commands to view, create, and navigate files and folders. We then focused on commands that help you administer and diagnose the network. We also covered commands that help you manage disks and filesystems. We then explored the most important aspect of the command line: getting help.

We concluded the chapter by covering the various network models that Windows is deployed in. You learned about various authentication methods, how to access resources, and how to configure the built-in firewall to allow applications to be accessed via the network.

Exam Essentials

Know the Windows installation process and options. You should know the various options when installing Windows 11, including the various partitions created and the differences between BIOS and UEFI firmware. You should also be familiar with the Windows 11 installation process.

Know which types of installations are possible with Windows. You should know which operating systems can be upgraded to Windows 11 and which require a clean installation.

Understand upgrading to Windows. You should know that a custom installation either wipes the old system or replaces the existing system, putting the old files into `WINDOWS.OLD`. Applications then have to be reinstalled, and user data has to be migrated from the old system using tools such as USMT. An upgrade preserves the existing applications and the user data, moving them into the new operating system.

Understand what each of the command-line utilities does. Many utilities allow you to perform various functions at the command line. These include `cd`, `dir`, `ipconfig`, `ping`, `netstat`, `nslookup`, `net use`, `tracert`, `pathping`, `chkdsk`, `format`, `diskpart`, `md`, `rmdir`, `robocopy`, `hostname`, `net user`, `winver`, `whoami`, `[command name] /?`, `gpupdate`, `gpresult`, and `sfc`. The `cmd` command opens a command prompt, where you can type the rest of the commands. If you're not sure how to use a particular utility, using the `/?` switch at the end of the command will provide information about how to use it.

Understand networking in Windows. You should understand the different networking models in which a Windows operating system can participate: workgroups, domains, and Microsoft Entra ID. You should also understand the basic user authentication methods used with each of these networking models, as well as common methods for accessing networking resources.

Know the various ways to establish a connection. You should know how to connect and diagnose wired and wireless connections to a network. In addition, you should understand how a WWAN is configured and how VPN networks are used and configured.

Know the various ways to access resources. You should know how to connect to network shares and printer shares using the UNC path of the resource. In addition, you should know the other types of resources, such as administrative shares and other hardware-based resources that you may encounter.

Know the basic elements of client configuration. You should be familiar with the Internet Protocol (IP) addressing scheme, Domain Name System (DNS) settings, subnet mask, gateway, and static vs. dynamic IP addressing.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. You are tasked with installing Windows 11 on a new computer. The computer has a 1 GHz processor, 2 GB of RAM, and 32 GB of storage. What should you do?
 - A. Proceed with the installation.
 - B. Upgrade the RAM to at least 4 GB.
 - C. Increase the storage to 64 GB.
 - D. Both B and C.
2. What is the name of the pass that detects hardware and installs drivers?
 - A. Generalize
 - B. OOBE
 - C. Specialization
 - D. WinPE
3. Which partition contains the recovery utility in the event of a problem?
 - A. BCD
 - B. System
 - C. ESP
 - D. WinRE
4. Which tool is used to create an operating system image for deployment?
 - A. WDS
 - B. MAP toolkit
 - C. USMT
 - D. sysprep
5. Which command will allow copying of all data to include NTFS permissions?
 - A. xcopy
 - B. copy
 - C. chkdsk
 - D. robocopy
6. Which Windows Update branch for Windows will install features as they are released to the general public?
 - A. Semi-Annual Channel
 - B. General Availability
 - C. Long-Term Servicing Channel
 - D. Insider Program

7. Which edition of Windows 11 will not allow for Insider Program branch releases?
 - A. Home
 - B. Pro
 - C. Education
 - D. Enterprise
8. You are configuring a new Windows 11 installation and need to ensure it supports Secure Boot. What type of firmware must the computer have?
 - A. BIOS
 - B. UEFI
 - C. Legacy BIOS
 - D. EFI
9. Which command is used to measure packet loss as a packet travels to a destination address?
 - A. ping
 - B. nslookup
 - C. pathping
 - D. tracert
10. Which tool allows you to report a remote computer's inventory of hardware?
 - A. regedit.exe
 - B. msinfo32.exe
 - C. msconfig.exe
 - D. dxdiag.exe
11. Which command will allow you to check a volume for corruption?
 - A. diskpart
 - B. format
 - C. chkdsk
 - D. sfc
12. Which restriction will be imposed on Windows users until they activate the operating system?
 - A. They won't be able to launch applications.
 - B. They won't be able to run Windows Updates.
 - C. They won't be able to change the wallpaper.
 - D. They won't be able to browse the Internet.

13. What is the name of the pass that configures and creates the user environment during the setup process?
 - A. Generalization
 - B. OOBE
 - C. Specialization
 - D. WinPE
14. When installing Windows, you can control telemetry data. In which step during setup is this configured?
 - A. Privacy settings
 - B. Locale options
 - C. Partitioning options
 - D. Account creation
15. You are troubleshooting a network issue and need to verify the path packets take to a destination. Which command should you use?
 - A. ping
 - B. tracert
 - C. nslookup
 - D. netstat
16. You want to upgrade from 32-bit Windows 10 Professional to a 64-bit version of Windows 11. What are your options?
 - A. Upgrade directly to Windows 11 Pro 64-bit.
 - B. Upgrade first to Windows 11 Pro 32-bit, and then upgrade to 64-bit.
 - C. Upgrade to Windows 10 64-bit, and then upgrade to Windows 11 64-bit.
 - D. You must perform a clean installation of Windows 11 Pro 64-bit.
17. Where is the boot configuration data stored on a Windows installation using EFI?
 - A. EFI System Partition
 - B. WinRE partition
 - C. Secure Boot partition
 - D. C:\WINDOWS
18. Which command is used to identify ports in use by applications and the operating system?
 - A. netstat
 - B. ipconfig
 - C. pathping
 - D. nslookup

19. Which type of connection is configured in the Cellular screen in the Settings app?
- A. Wired
 - B. Wireless
 - C. WWAN
 - D. VPN
20. What is the maximum number of concurrent connections that can be made to a Windows workstation?
- A. 30 connections
 - B. 40 connections
 - C. 20 connections
 - D. 25 connections

Performance-Based Question

You will encounter performance-based questions on the A+ exam. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exam. To see how your answers compare to the authors', refer to Appendix B.

You need to join a Windows 11 workstation to an Active Directory domain. What are the steps you need to follow to complete the task?

Chapter 4

Working with macOS and Linux

**THE FOLLOWING COMPTIA A+ 220-1202
EXAM OBJECTIVES ARE COVERED IN THIS
CHAPTER:**

- ✓ 1.0 Operating Systems
- ✓ 1.8 Explain common features and tools of the macOS/
desktop operating system.
 - Installation and uninstallation of applications
 - File types
 - .dmg
 - .pkg
 - .app
 - App Store
 - Uninstallation process
 - System folders
 - /Applications
 - /Users
 - /Library
 - /System
 - /Users/Library
 - Apple ID and corporate restrictions
 - Best practices
 - Backups
 - Antivirus
 - Updates/patches
 - Rapid Security Response (RSR)
 - System Settings



- Displays
- Networks
- Printers
- Scanners
- Privacy
- Accessibility
- Time Machine
- Features
 - Multiple desktops
 - Mission Control
 - Keychain
 - Spotlight
 - iCloud
 - iMessage
 - FaceTime
 - Drive
 - Gestures
 - Finder
 - Dock
 - Continuity
- Disk Utility
- FileVault
- Terminal
- Force Quit

✓ **1.9 Identify common features and tools of the Linux client/desktop operating system.**

- File management
 - ls
 - pwd
 - mv



- cp
- rm
- chmod
- chown
- grep
- find
- Filesystem management
 - fsck
 - mount
- Administrative
 - su
 - sudo
- Package management
 - apt
 - dnf
- Network
 - ip
 - ping
 - curl
 - dig
 - traceroute
- Informational
 - man
 - cat
 - top
 - ps
 - du
 - df
- Text editors
 - nano



- Common configuration files
 - `/etc/passwd`
 - `/etc/shadow`
 - `/etc/hosts`
 - `/etc/fstab`
 - `/etc/resolv.conf`
- OS components
 - `systemd`
 - `kernel`
 - `bootloader`
- Root account



CompTIA has acknowledged that system administrators and technicians are increasingly dealing with more than just Windows on a daily basis. Therefore, they have included exam objectives based on *macOS* and *Linux*.

This chapter looks at the non-Windows operating systems from the standpoint of what you need to know to pass the exam. All the topics relevant to Exam Objectives 1.8 and 1.9 of the 220-1202 exam are covered.

macOS and Linux

In the beginning there was UNIX. UNIX System 5 (version 5) is an operating system originally created and licensed by AT&T Labs. The UNIX operating system is considered to be the root of all UNIX-based operating systems. In the mid-1970s, the University of California at Berkeley (UC Berkeley) licensed UNIX from AT&T for their computer systems and expanded on the tools shipped in the original version of UNIX. These tools became the foundation of UNIX as it is today, but UC Berkeley only licensed the operating system for specific machines. The students quickly became upset and developed and released a version of UNIX called the Berkeley Software Distribution (BSD). The term *distribution* is used today with UNIX/Linux operating systems to define the operating system and its ecosystem for application management, patching, and upgrades.

Although it began with UNIX, BSD became very popular because it was an open source license. This allowed everyone to use the operating system on any computer system they wished. In the mid-1980s, Steve Jobs created a company called NeXT and built computers that furnished the NeXTSTEP operating system. The NeXTSTEP operating system was originally built from BSD version 4.3. Unfortunately, the NeXT computer company never really took off. However, the NeXTSTEP operating system was acquired by Apple and eventually became the macOS we know today.

Linux has a very different origin story from macOS. Actually, Linux has nothing to do with the original codebase of UNIX. In the mid-1990s a Finnish student named Linus Torvalds set out to create a completely open source operating system for the world to use. Linux was the result of his efforts; it was designed from scratch, so it was completely free for anyone to use or incorporate into their own products. Today you can find a great number of Linux distributions, such as *Ubuntu*, *Debian*, Arch Linux, Gentoo, *Red Hat* . . . and the list goes on.

Although Linux has a completely different codebase from BSD and UNIX, the operating system itself functions similarly. Only the kernel and interworkings of the OS are different.

Many of the applications that were created by students on the BSD platform were ported over to Linux. Functionally, the operating systems are very similar in design and usability.



The complete history of UNIX is vast and interesting; we have covered only the big events in the UNIX timeline. Many different organizations were involved, spanning all the way back to the 1960s. Learn more about the history of UNIX here: https://UNIX.org/what_is_UNIX/history_timeline.html.

Applications on macOS

The applications that are available or installed on an operating system indirectly define an OS by extending functionality to the end user. As with the case of UNIX, BSD, Linux, and macOS, all of the basic command-line applications are similar in functionality, as you will learn in this chapter. These command-line applications are preinstalled with the OS distribution. When you need functionality that isn't part of the base OS, you have a few different ways to install applications.




Exam Objective 1.8, “Explain common features and tools of the macOS/desktop operating system,” requires you to know about installing applications on macOS, system folders, Apple ID, best practices, system preferences, and features of macOS.

Installing Applications from the App Store

macOS contains an application ecosystem called the *App Store*, shown in Figure 4.1. From the App Store you can download and buy applications for macOS. Some applications might be free from the App Store, but they contain in-app purchases. They are called *freemium* games and applications, and they typically use this feature.

Applications installed through the App Store do not function any differently from legacy downloaded applications. Unfortunately, not all applications are available through the App Store. The developer of an application must publish it on the App Store, and not every vendor will submit their application. This is mainly due to the costs related to publishing an app to the App Store.

The benefit of using the App Store to install applications for macOS is if you purchase an app and have two devices, then you need to purchase the application only once. This is assuming you are logged in with the same Apple ID that purchased the application. Also, updates for the applications are automatically installed, compared to downloading applications from the developer directly. The Updates section in the App Store will show you all the applications containing updates.

FIGURE 4.1 macOS App Store

When you first turn on your Apple device and run through the installation, it will ask for an *Apple ID*. You really can't move on with the setup without an Apple ID. The Apple ID is your digital identity on the device, and it's what also ties your Apple Wallet to the App Store for purchases. The Apple ID will contain credit card information that can be used for purchases.

An organization may use mobile device management (MDM) software to control the installation of applications from the App Store or any software installations in general. When an organization uses MDM to restrict the installation of applications, this is considered corporate restrictions. An organization might also use the App Store as a distribution point for applications that its employees are expected to use. The organization's MDM software is associated with the employee's Apple ID, which is usually the employee's email address. These features were introduced in macOS starting with version 10.9. New features are introduced in every version of macOS.

Installing Downloadable Applications

Installing downloadable applications on macOS is not much different from the same process in Windows. The process consists of three steps: providing the installation to the operating system, mounting the installation, and then installing the applications.

The first step requires you to provide the application's installation files to the operating system. This step is normally done nowadays as a download, since we always want the latest updated applications. The process is pretty routine, as you can imagine; you navigate to a

web page and download the app. However, providing the application might also be as simple as plugging in a USB drive. macOS will automatically mount the USB drive onto the desktop and open the root folder.

The second and third step in the installation process is to launch the file you downloaded and install the application. When you launch the file, what will happen next depends on the type of file. There are a number of different file types that can be downloaded from the Internet, such as ZIP, ISO, DMG, PKG, or APP, just to name a few. The most common files for macOS applications are as follows:

APP Files The APP filename will end in `.app` and it is used to contain the application. The APP file is actually a container for various resource files and folders, and it also contains the executable for launching the application. You will rarely download an APP directly from the Internet. The APP file will often be found inside a ZIP, DMG, or ISO file or physical media.

DMG Files A DMG filename will end in `.dmg` and it stands for disk imaging file. The DMG file is directly comparable to ISO or ZIP files. The DMG file contains a file structure, and the file can be mounted as a disk on the macOS desktop. DMG files used for application installation will normally contain two files: the application itself and a shortcut to the Applications folder. This allows you to drag (copy) the APP file into the Applications folder, and doing so will install the application.

PKG Files A PKG filename will end in `.pkg` and it is an automated package installer. The PKG file is comparable to the application installation programs that are commonly found in Windows. When the PKG file is launched, it will ask the user a series of questions about the installation. PKG files are found inside DMG, ZIP, and physical media. These files are used when the installation requires customization by the end user.



Exam Objective 1.8 specifically calls out DMG, PKG, and APP file types. Be sure you are familiar with them for installation of applications in macOS.

In Exercise 4.1, you will install an application on macOS to understand the process.

EXERCISE 4.1

Installing Applications on macOS

1. Open the Safari app from the Dock.
2. When Safari opens, type **sublimetext.com** in the address bar and press Return.
3. When the web page loads, choose the download option of macOS.
4. A prompt will ask you to allow downloads from `www.sublimetext.com`; click Allow.
5. The download is stored in the Downloads folder located on the right side of the Dock. Click the Downloads folder and select the application `Sublime Text Build 4189.dmg`.

6. After macOS mounts the DMG as a drive and opens the folder, drag the SublimeText icon over the Applications shortcut and release it.
7. You can then eject the SublimeText mount on the desktop by selecting the folder and clicking Finder > File > Eject.
8. Press Shift+Command+A to open the Applications folder.


Managing Applications

The Applications folder contains all the applications installed on macOS, as shown in Figure 4.2. You can view and manage applications inside this folder. The two most common methods of launching the folder are using the *Finder app* and selecting Go > Applications or pressing Shift+Command+A. This key sequence will also launch the Applications folder and requires no shuffling of applications to get to the Finder app.

From the Applications folder, you can delete user-installed applications. However, preinstalled applications in this folder cannot be deleted by the user, since they are technically part of the operating system. These applications are similar to the built-in applications on Windows.

To delete a user-installed application, select the folder; then on the Finder menu, select File > Move To Trash. This will place the application into the Trash and remove it from the operating system. Another method is to select the application and press Command+Delete. In Exercise 4.2 you will uninstall an application to understand the process better.

FIGURE 4.2 macOS Applications folder



EXERCISE 4.2**Uninstalling Applications on macOS**

1. Press Shift+Command+A to open the Applications folder.
 2. Select the SublimeText application.
 3. On the Finder menu, click File ➤ Move To Trash.
 4. Open the Trash folder from the right side of the Dock and observe that the Sublime Text application is there.
-

Creating Shortcuts

The last and most important topic we will cover about applications is the availability of the application. Creating shortcuts is an essential part of your overall workflow. You should be able to launch an application in as few clicks as possible. To create a shortcut, launch the Applications folder by pressing Shift+Command+A and drag the application to the *Dock*. The dock is covered later in this chapter in the “macOS Components” section.

macOS Components

macOS is renowned for its robust system utilities and comprehensive settings that empower users to optimize their computing experience. From essential tools like Disk Utility, which manages and repairs your storage devices, to the Accessibility settings that ensure inclusivity for all users, macOS offers a wide array of features designed to enhance productivity and usability. The System Settings provide a centralized hub for customizing your Mac, allowing you to tailor everything from privacy controls to display preferences. Additionally, features like Continuity and iCloud integration ensure a seamless experience across all your Apple devices. In the following sections, we’ll delve deeper into these utilities and settings, exploring how they can help you get the most out of your macOS environment.

Interestingly, macOS shares several similarities with Windows 11, making it easier for users familiar with one operating system to transition to the other. Both operating systems offer centralized settings menus—System Settings in macOS and Settings in Windows 11—where users can manage system preferences, update software, and customize their experience. Additionally, both macOS and Windows 11 provide robust accessibility features, ensuring that users with diverse needs can effectively use their devices. The integration of cloud services, such as iCloud for macOS and OneDrive for Windows 11, further enhances the user experience by enabling seamless access to files and data across multiple devices. This familiarity in structure and functionality helps users navigate and utilize the features of both operating systems with ease.

Tools

There are a number of tools to be aware of in macOS and Linux. Most of these have counterparts in the Windows world, and we'll make comparisons where they apply. Tools are released on a daily basis, but the following are the most important for the CompTIA exam objectives and daily maintenance:

Restore/Snapshot *Snapshots* are archives of key files and settings as they exist at a moment in time. It is possible to take snapshots and restore them as needed. The macOS Time Machine feature performs a backup of your data to the external drive connected to the system. However, the external drive may not always be connected to the system, so local snapshots are also created. These local snapshots allow you to restore files using the local filesystem. To use local snapshots, start Time Machine, and if the external storage device is not connected, you'll be presented with the local snapshots. Snapshots in macOS are similar to Windows 11 System Restore points.

Image Recovery As a general rule, images are typically larger than snapshots. You can take a snapshot of a project, and that will include all the files associated with the project, whereas an image would include the project files and all files on the system at the time. Again, this is only a general rule, since images can be granular as well. Typically, however, snapshots are thought of as subsets of images.

The macOS *Disk Utility* can be used to create an image of the macOS operating system, and the image can be directed to an external storage device. Linux can use a multitude of open source tools to create an image of the operating system. The most common is the `dd` command.

Disk Maintenance Utilities A number of disk maintenance utilities are available with, or for, macOS and Linux. The most important utilities you should know are `du`, which shows how much disk space is in use; `df`, which shows how much space is free; and `fsck`, which checks and repairs disks. Although these command-line tools are available on both macOS and Linux, Disk Utility is available only on macOS. The Linux operating system can use a variety of open source disk utilities.

Disk Utility on macOS is a versatile tool designed for managing your disks and storage devices. It allows you to perform a variety of tasks such as formatting drives, creating disk images, and repairing disk errors. You can also use Disk Utility to partition your drives, and securely erase data. To launch Disk Utility, open Finder > Applications > Utilities > Disk Utility. This powerful tool is essential for maintaining the health and efficiency of your Mac's storage system. The macOS Disk Utility is shown in Figure 4.3. This utility is comparable with the Windows Disk Management MMC.

FileVault The *FileVault* tool allows you to provide full-disk encryption to prevent unauthorized access to the information contained on the device during startup. FileVault uses XTS-AES-128 with a 256-bit key. The feature first appeared in OS X 10.3 (2003), and it is currently on its second release (FileVault 2). It can be enabled by navigating to

System Settings ➤ Security & Privacy ➤ FileVault, and from there you can click Turn On FileVault, as shown in Figure 4.4. This utility is comparable to the Windows 11 BitLocker functionality.

FIGURE 4.3 macOS Disk Utility





FIGURE 4.4 FileVault encryption



Shell/Terminal The command-line interface in Linux is usually reached through a Terminal session (though that need not always be the case). It is in this Terminal session that you interact with the *shell*—the interpreter between the user and operating system. The most popular shell for Linux today is *Bash* (*an acronym for Bourne Again Shell*), but *csh* (C-shell), *ksh* (Korn shell), and a number of others are also in use. In the macOS/iOS world, *OpenSSH* (*open shell*) is often downloaded and installed. The macOS Terminal utility is accessible by going to Finder > Applications > Utilities > Terminal. Windows 11 has also adopted the concept of a terminal for shell access, this feature is comparable to the Windows Terminal.

Screen Sharing This feature, built into macOS, enables you to share your screen with others. While your screen is shared, the user of the other computer sees what you see, and they can open, move, and close files as well as restart your system. It requires a *Virtual Network Computing (VNC)* connection client.

**NOTE**


The Internet offers numerous versions of VNC connection clients—many are free, and some have unique features that you must pay for. An example of a free VNC connection client is TightVNC (www.tightvnc.com/download.php).

Force Quit Sometimes an application can become unresponsive. When that happens in macOS, you can use *Force Quit* to force the application to close. The Force Quit dialog box can be accessed by pressing Option+Command+Esc. This is similar to pressing Ctrl+Alt+Delete on Windows. Alternatively, you can choose Force Quit from the Apple menu in the upper-left corner of the screen. Most devices, whether running macOS or iOS, offer similar options. With an iPad, for example, if you press the Home button twice quickly, a list of apps will appear at the bottom of the screen. Swipe to get to the app in question, and then swipe up on its preview to close it.

System Settings

Like any other modern-day operating system, macOS is highly customizable. macOS has a feature called System Settings that allows you to customize the operating system, similar to the Windows 11 Settings app and Control Panel applets. In this section, we will cover the most important System Settings. After reading this section, you'll understand how to configure macOS and personalize it for your needs.

All settings are accessed through the System Settings screen, shown in Figure 4.5. The System Settings screen replaced the System Preferences screen starting in macOS Ventura 13.7.2. The System Settings screen can be accessed in a number of ways. The easiest way to launch it is from the Dock, but you can also launch it by clicking the Apple icon in the upper left of the desktop and selecting System Settings.


FIGURE 4.5 macOS General System Settings

The number of icons on the System Settings screen depends on the applications installed on the operating system and if they are configurable. Let's explore the various System Settings that you need to know for the CompTIA exam:

Displays The settings on the Displays tab are primarily used to adjust display resolution and brightness, as shown in Figure 4.6. The Displays tab will also allow you to adjust the operating system if you have multiple monitors connected. You can also


adjust the color tones the display produces by selecting the Color Profile and calibrating the display profile. This setting is generally used when editing graphic artwork is the primary function of the device. The Night Shift option is used to adjust colors at night so that you can get a better night's sleep. This feature is primarily used with laptops, since, for example, you may check the news or catch up on emails before you fall asleep. The device will automatically make the system colors warmer so that it is easier to fall asleep afterward.

FIGURE 4.6 Displays tab



Network The Network tab contains settings pertaining to the network connectivity for the device, as shown in Figure 4.7. In the figure, Wi-Fi (wireless) is the main connectivity method. If the device used a wired connection, the connectivity method would be Ethernet. You can open the Network settings by clicking the Apple icon in the upper left of the desktop and selecting System Settings and then Network.

FIGURE 4.7 Network tab




Just like Windows 11, macOS automatically detects locations, and you can change your Network setting based on these locations. For example, you could have a location of work and home. At work you might turn off Bluetooth, and at home Bluetooth might be on. This is achieved by clicking on the three dots with the down arrow under Other Services and choosing Locations ➤ Edit Locations; by default, it is set to Automatic.

The Network tab is also where you can join wireless networks and change how you join wireless networks. By default, the device will automatically join the network selected. However, there are circumstances where you would not want to automatically join the network, such as if you were directly connecting to a wireless device like a camera that broadcasts its own Service Set Identifier (SSID). If the primary wireless network was still set to automatically join, the device would keep disconnecting from the camera to join the primary network.

If you click Details, you can change your primary wireless networks and specify whether they are auto-joined, as shown in Figure 4.8. In addition to changing advanced properties for the wireless connection, you can select the TCP/IP tab and statically set the IP address. The DNS tab allows you to change the DNS servers to be queried. The WINS tab is for a deprecated Windows service, Windows Internet Name Service, that permits network browsing via broadcasts. The 802.1X tab allows you to set up 802.1X profiles for network-level security. The Proxies tab allows you to configure proxy servers for traffic on the device, as well as bypass local and select addresses. Finally, the Hardware tab allows you to set special characteristics based on the device.


FIGURE 4.8 Details Network settings



Each device will have various methods for network connectivity. Depending on the various methods, the Network settings will differ based on the connection. The Details settings may also differ depending on the type of connectivity method for the device.


The Firewall option allows you to turn on and configure the built-in firewall for macOS. It is not turned on by default, but you can turn it on by clicking Turn On Firewall and then configure the firewall options, as shown in Figure 4.9. You can choose to block all inbound connections by default and create exceptions for only the applications you choose. By default, after turning on the firewall, the operating system allows all inbound connections to applications running on the system. You must choose to block all inbound connections and configure the exceptions.

FIGURE 4.9 Firewall settings



Printers & Scanners The Printers & Scanners tab, shown in Figure 4.10, is where you can access all settings related to printing and scanning for the device. In the figure, only an HP DeskJet Plus 4100 is connected to the device. However, there could be several different printers simultaneously connected via this System Setting. This would then allow the user to select any of the printers for their use in the desired application.

FIGURE 4.10 Printers & Scanners tab



Printers can be added and removed with this System Setting by clicking the Add Printer, Scanner, or Fax on the Printers & Scanners screen. You can then select the device and open the print queue by clicking Print Queue. This will allow you to see all the print jobs currently waiting to be printed from the local device. You can also click Options & Supplies and view the various options for the printer and check the ink or toner supply levels. In addition, you can share the printer by clicking Open Sharing Settings. You can then configure the Sharing settings and choose who can print to the printer.

Many printers today can also be purchased as multifunction copiers (MFCs), also called multifunction devices (MFDs) or multifunction printers (MFPs), which means the printer doesn't just print—it can copy and scan as well. If the device is capable of scanning, a Scan tab will be available after you select the device. Depending on the MFC device attached, the Scan tab will allow you to configure various settings. Although MFC devices are becoming common, if a stand-alone scanner was connected to the device, it would show in the Printers & Scanners settings. The device would have only one tab for scanning.

In addition to changing settings for the printers and scanners attached, you can change the default printer and the default paper size. The default printer is set to the last printer used.

Privacy & Security The Privacy & Security setting contains a number of settings that apply to the privacy of the user and the security of the device. The first setting you will notice is Location Services, as shown in Figure 4.11. Location Services on macOS provide a seamless way to enhance your user experience by allowing apps and websites to gather and use information based on your current location. You can also control the applications that request access to your data, such as Contacts, Calendars, Reminders, and Photos, just to name a few. You can add or remove the applications that need or have access to your data.


Just like in Windows 11, you can view all of the applications installed on the macOS operating system and their given permissions. Users can review and adjust permissions for each app, ensuring that only trusted applications can access their data. This granular control helps protect personal information from unauthorized access, providing peace of mind and a secure computing environment.

The Security section of the Security & Privacy screen contains some security controls. The FileVault option allows you to configure disk-level encryption to protect your files in the event the device is lost. You can turn on the FileVault feature by clicking Turn On FileVault. Once you do, you will need to unlock the disk with the user's password. If there are multiple users configured on the device, they will need to verify their passwords before the encryption is performed.

Lockdown Mode in macOS, is a feature also found in the Security section of the Privacy & Security screen. It offers an extreme level of protection for users who might be targeted by sophisticated cyberthreats. When enabled, Lockdown Mode significantly reduces the attack surface by limiting the functionality of certain apps, websites, and features. This includes blocking most message attachments, disabling complex web technologies, and restricting incoming FaceTime calls to known contacts. While this

mode enhances security, it also impacts the overall functionality and performance of the device, making it a trade-off between usability and protection. Lockdown Mode is designed for those who require the highest level of security, ensuring their sensitive data remains safeguarded against advanced digital threats.


FIGURE 4.11 Privacy & Security



Accessibility The Accessibility tab allows you to customize macOS to support requirements such as vision, hearing, motor skills, or speech, as shown in Figure 4.12. In the Vision section, you can turn on VoiceOver, which will provide spoken commands as well as descriptions of items in Braille. Turn on Zoom to allow zooming with the use

of the keypad. The Display settings allow you to turn on display features for high contrast as well as reduce motion. The Spoken Content section allows you to configure the operating system to speak announcements, sections, items, and typing feedback. Use the Audio Descriptions section to turn on additional spoken content.

FIGURE 4.12 Accessibility System Settings



The Hearing section is designed to support users with hearing impairments. It offers features like mono audio, which combines stereo sound into a single channel, making it easier for users who are deaf or hard of hearing in one ear. The section also allows you to control hearing devices, as well as visually indicate when audio alerts are being

played so that the user does not miss an alert. The Captions option will display captions (subtitles) for the operating system. The Live Captions option will allow you to control how apps, like FaceTime, perform live captioning.


The Motor section is tailored to assist users with motor impairments. The Voice Control option provides the ability to speak to the computer to display and edit text. The Keyboard option lets you enable sticky keys and slow keys. The Pointer Control option lets you control how the trackpad operates by changing click speed. The Switch Control option allows you to specify that an adaptive device (such as a joystick) be used to control your Mac, enter text, and interact with items on your screen.

The Speech section is designed to assist users with speech impairments. It includes features like Live Speech, which allows users to use Text to Speech for applications like FaceTime. The Personal Voice option allows you to custom tailor a voice that sounds like yourself. The Vocal Shortcuts allows you to teach the operating system to recognize a custom phrase that performs an action inside the operating system.

The Siri section allows you to configure Siri to accept typed requests in lieu of spoken requests. In the Shortcut section, you can specify that a shortcut list appear when you press Option+Command+F5.

Backup/Time Machine As discussed in the previous section, backups are important. In the macOS world, *Time Machine* is a feature that can be used to make backups of various types (incremental, full, and so forth), as shown in Figure 4.13. In order to use Time Machine to back up your data, you will need an external storage device. These devices can be connected via external USB, Thunderbolt, or wirelessly with 802.11. To access Time Machine, select Apple menu > System Settings > General > Time Machine. When an external storage device is connected for the first time and Time Machine is not configured, macOS will prompt you to set up Time Machine.

FIGURE 4.13 Time Machine settings




Features



Exam Objective 1.8, “Explain common features and tools of the macOS/ desktop operating system,” requires you to be familiar with multiple desktops, Mission Control, Keychain, Spotlight, iCloud, iMessage, FaceTime, Drive, Gestures, Finder, Dock, Continuity, Disk Utility, FileVault, Terminal, and Force Quit.

Multiple Desktops/Mission Control It is possible to run a large number of things at one time, whether those things are apps or desktops. Apple’s *Mission Control* is an easy way to see what is open and switch between applications. To access Mission Control, press the Mission Control key on an Apple keyboard, click the Mission Control icon on the Dock (or Launchpad), or swipe up with three or four fingers on a trackpad. Once Mission Control is accessed, all open windows and spaces are shown—grouped by app—and you can choose between them. The Apple macOS Mission Control feature is identical in functionality to Microsoft’s Task View feature. It allows you to see all the applications running on one screen so that you can choose the foreground application, as shown in Figure 4.14. By clicking the + symbol at the upper-right corner of the screen you can create a new virtual desktop.


FIGURE 4.14 Viewing multiple apps with Mission Control




You can access Mission Control in a few different ways. The first is with trackpad gestures by swiping up with 3 or 4 fingers on your trackpad. The second is the key strokes of Control+Up arrow. The third way is the Mission Control (F3) button on the keyboard. The last method to initiate it is the Mission Control icon in the dock.

Keychain Access/Passwords *Keychain Access* is a password management system from Apple. It was originally introduced in Mac OS X, and remained in macOS for some time. However, with the introduction of macOS 15.2 (Sequoia) it has been migrated to the Passwords app. The Passwords app allows you to store passwords for websites, mail servers, Wi-Fi, and so forth, as shown in Figure 4.15. There is an iCloud variant (iCloud Passwords) that keeps such information as Safari usernames/passwords and credit card information. Values stored in the Passwords app are kept current (changes are synced) to simplify resource access. To view the current Passwords, select Applications > Passwords.

FIGURE 4.15 Apple Passwords app



Spotlight The search feature within macOS is Spotlight, and a magnifying glass icon in the upper-right corner of the menu bar represents it (or you can press Command+spacebar from any app). As you type in the Spotlight utility, the results will display. As shown in Figure 4.16, we have searched for *apple*, and one of the results is the contact info for Apple Inc. Spotlight can search for documents, images, apps, and so on. Recent versions include autocomplete features as well as suggestions for additional results (usually on Wikipedia, *iTunes*, and such).

FIGURE 4.16 Apple Spotlight utility

iCloud Apple iCloud is a comprehensive cloud storage and cloud computing service that seamlessly integrates with Apple's ecosystem, providing users with a unified platform to store, sync, and access their data across all their Apple devices. iCloud ensures that your photos, documents, apps, and other important information are always up-to-date and accessible, no matter which device you're using. This service not only enhances productivity but also offers robust security features to protect your data. The iCloud settings can be seen in Figure 4.17.

One of the key components of iCloud is iCloud Drive, which allows users to store files and documents in the cloud and access them from any device. Whether you're working on a Mac, iPhone, or iPad, iCloud Drive ensures that your files are always within reach. You can organize your files into folders, share them with others, and even collaborate in real-time on documents. This makes iCloud Drive an essential tool for both personal and professional use.

iMessage is another integral part of iCloud, providing a seamless messaging experience across all Apple devices. With iMessage, your conversations are synced in real time, allowing you to start a chat on your iPhone and continue it on your Mac or iPad

without missing a beat. iCloud ensures that your messages are securely stored and backed up, so you never lose important conversations. Additionally, iMessage supports rich media, such as photos, videos, and even animated effects, making your communication more engaging.

FIGURE 4.17 iCloud configuration settings on macOS



FaceTime, Apple's video and audio calling service, is also tightly integrated with iCloud. FaceTime allows you to make high-quality video and audio calls to other Apple users, regardless of which device they are using. With iCloud, your FaceTime call history is synced across all your devices, making it easy to keep track of your communications.

Whether you're catching up with family, collaborating with colleagues, or having a virtual meeting, FaceTime provides a reliable and secure way to stay connected.

Together, these services—iCloud Drive, iMessage, and FaceTime—form an ecosystem under iCloud, enhancing the overall user experience by providing seamless integration, synchronization, and security across all Apple devices.



There can be costs associated with using iCloud. For example, the iCloud account on the macOS shown in Figure 4.17 can access 5 GB for free, but anything beyond that requires payment (\$0.99/month for 50 GB, \$2.99/month for 200 GB, or \$9.99/month for 2 TB, as of this writing). For the latest costs visit <https://support.apple.com/en-us/108047>.


Gestures With Apple products, it is possible to scroll, tap, pinch, and swipe to interact with the macOS or other products in a way that is intended to be natural and intuitive. You can accept the default actions for these gestures, or you can configure them differently, as shown in Figure 4.18. You can access them by clicking the Apple menu ➤ System Settings ➤ Trackpad.

To see the basics of gestures on a macOS, visit <https://support.apple.com/en-us/102482>.



Finder To use an analogy, the Finder is to macOS as File Explorer is to Windows. It lets you browse through folders and find files, disks, apps, and so on. For example, Figure 4.19 shows an example of the Finder displaying the contents of the *Utilities* folder. You can change the view to see the entries with images, in a list view, in a column view, and so on.

Dock In macOS, a bar of icons runs along the bottom (or side, if so configured) of your screen, as shown in Figure 4.20. That set of icons, known as the *Dock*, provides easy access to key apps that come with the Mac (such as Safari, Mail, App Store, and Music) or others that you choose to add there. To see the basics of the Dock on a Mac, visit <https://support.apple.com/guide/mac-help/open-apps-from-the-dock-mh35859/mac>.

Continuity Continuity is a suite of features in macOS designed to create a seamless experience across all your Apple devices. It allows you to start a task on one device and continue it on another without missing a beat. For example, with Handoff, you can begin writing an email on your iPhone and finish it on your Mac. Universal Clipboard lets you copy text, images, or videos on one device and paste them on another. Continuity Camera enables you to take a photo or scan a document with your iPhone and instantly insert it into a document on your Mac. To use Continuity, ensure all your devices are signed in to the same Apple ID, have Bluetooth and Wi-Fi enabled, and are within close proximity to each other.

FIGURE 4.18 Settings for default gestures on macOS

There are many other features of macOS that make it a powerful operating system to use. These features, however, are the ones that CompTIA wants you to be aware of for the exam. Make sure that you know the purpose of each as you prepare.

FIGURE 4.19 Apple Finder utility**FIGURE 4.20** Apple macOS Dock

macOS System Folders

The macOS operating system organizes its filesystem into several key folders, each serving a distinct purpose to ensure the smooth operation and user experience of the operating system. Understanding these folders is crucial for both everyday users and system administrators. Each folder has a specific role, from housing system-critical files to storing user data and application resources. This organization not only helps in maintaining system stability and performance but also aids in troubleshooting and system management.

Application developers also have specific expectations for these folders. They rely on the standardized structure to ensure their applications can find and store necessary files correctly.

For instance, developers expect to place their applications in the `/Applications` folder, store system-wide resources in `/Library`, and save user-specific settings in `/Users/username/Library`. This consistency allows applications to function seamlessly across different macOS installations, providing a reliable and predictable environment for both development and usage. By adhering to these conventions, developers can create applications that integrate smoothly with the macOS ecosystem, enhancing the overall user experience.



Exam Objective 1.8, “Explain common features and tools of the macOS/desktop operating system,” calls out `/Applications`, `/Users`, `/Library`, `/System`, and `/Users/Library` System folders in macOS. Be sure you know how to access and use them.

/Applications

The `/Applications` folder is a central hub for all the applications installed on a macOS system. This directory is designed to be accessible to all users, meaning that any application placed here can be used by anyone with an account on the macOS. It contains both system applications, which come pre-installed with macOS, and third-party applications, which users can install from the App Store or download from the Internet.

By default, when you install an application from the App Store or download one from the Internet, it is placed in the `/Applications` folder. This standard location makes it easy for users to find and launch their software, as they know exactly where to look for their applications. This centralized approach helps keep the system organized, as all applications are stored in one place rather than scattered across different directories.

Having a single location for applications also simplifies the process of managing and updating software. Users can easily see all installed applications in one view, making it straightforward to update, uninstall, or troubleshoot any application. For example, when a new version of an application is released, users can quickly locate the application in the `/Applications` folder and update it. Similarly, if an application needs to be removed, users can simply drag it from the `/Applications` folder to the Trash.

Moreover, this organization benefits system administrators and developers. System administrators can efficiently manage software installations across multiple user accounts, ensuring that all necessary applications are available to all users. Developers, on the other hand, can design their applications with the expectation that they will be installed in the `/Applications` folder, allowing them to streamline the installation process and ensure compatibility with macOS conventions.

/Users

The `/Users` folder is a fundamental part of the macOS filesystem, serving as the home directory for all user accounts on the system. Within this folder, each user has a dedicated subdirectory, often referred to as their home directory, which contains all their personal files,

settings, and preferences. This structure is crucial for maintaining an organized and secure environment, especially on a shared Mac.

Each user's home directory typically includes several standard folders such as Documents, Downloads, Desktop, and Pictures. These folders help users organize their files and make it easy to locate specific types of data. For example, the Documents folder is intended for storing text files, spreadsheets, and other document types, while the Pictures folder is for images and photos.

The separation of user data into individual home directories ensures that each user's information is kept private and secure. This means that one user cannot easily access another user's files, providing a layer of privacy and security. Additionally, this structure allows users to customize their settings and preferences without affecting others. For instance, each user can have their own desktop background, application settings, and browser bookmarks.

Moreover, this organization simplifies system management and troubleshooting. System administrators can manage user accounts and data more efficiently, knowing that each user's files are contained within their home directory. It also makes it easier to back up and restore user data, as each home directory can be treated as a separate unit.

/Library

The `/Library` folder contains systemwide resources and settings used by macOS and installed applications. This includes fonts, application support files, and system preferences. There are actually three Library folders in macOS: the system Library (`/System/Library`), the local Library (`/Library`), and the user Library (`/Users/username/Library`). Each serves a different scope, from systemwide resources to user-specific settings. The `/Library` folder is crucial for the proper functioning of applications and the operating system, as it stores essential files that applications rely on to run correctly.

System Library (`/System/Library`) The System library, located at `/System/Library`, contains the core components and resources that macOS needs to operate. This includes essential system frameworks, system extensions, and default fonts. These files are integral to the operating system's functionality and are typically protected to prevent accidental modification. Applications often rely on the frameworks and libraries found in this directory to perform system-level tasks and ensure compatibility with macOS.

Local Library (`/Library`) The local library, found at `Library`, stores resources and settings that are shared across all users on the Mac. This includes application support files, systemwide preferences, and shared libraries. For example, when an application is installed, it might place shared resources such as plug-ins, configuration files, and documentation in this directory. This ensures that these resources are available to all users, providing a consistent experience regardless of who is logged in. The local library also contains items like fonts and audio plug-ins that can be used by any application on the system.

User Library (`/Users/username/Library`) The user library, located at `/Users/username/Library`, is specific to each user account. This directory contains

user-specific application support files, caches, preferences, and other data. It is hidden by default to prevent accidental modification, but it can be accessed by advanced users who need to troubleshoot or customize their applications. Each user has their own Library folder within their home directory, ensuring that their settings and data are kept separate from other users on the system. This allows applications to store personalized settings and data, providing a tailored experience for each user.

By organizing resources and settings into these three distinct Library folders, macOS ensures that systemwide and user-specific data are properly managed and isolated. This structure helps maintain system stability, security, and performance, while also providing flexibility for both users and developers. Applications can efficiently access the resources they need, whether they are shared across the system or specific to an individual user, ensuring a seamless and reliable user experience.

/System

The /System folder is a critical directory that houses the core components of macOS. This includes the operating system kernel, essential system frameworks, and other low-level files necessary for macOS to function. This folder is typically read-only to prevent users from making changes that could destabilize the system. The /System folder ensures that the fundamental parts of the operating system remain intact and secure, providing a stable environment for the user.

The /System folder is meticulously organized to include various subdirectories that each play a vital role in the functioning of macOS. One of the key components within this folder is the operating system kernel, which is the core part of macOS responsible for managing system resources, hardware communication, and overall system stability. The kernel acts as a bridge between the hardware and software, ensuring that applications can run smoothly and efficiently. The following are some of the important functions that the /System folder provides:

System Frameworks Another crucial element within the /System folder is the system frameworks. These frameworks provide a set of standardized interfaces and libraries that developers use to build applications. By relying on these frameworks, applications can perform complex tasks such as graphics rendering, data processing, and network communication without needing to implement these functionalities from scratch. This not only simplifies the development process but also ensures that applications are compatible with the macOS environment.

System Extensions and Drivers The /System folder also contains system extensions and drivers that enable macOS to interact with various hardware components. These extensions are essential for the proper functioning of peripherals such as printers, scanners, and external storage devices. By keeping these files within the /System folder, macOS ensures that they are protected from accidental deletion or modification, which could otherwise lead to hardware malfunctions or system instability.

Default System Applications Additionally, the `/System` folder includes default system applications and utilities that are integral to the macOS experience. These applications, such as the Finder, Terminal, and Activity Monitor, provide essential functionalities that users rely on for everyday tasks. By housing these applications within the `/System` folder, macOS guarantees that they are always available and functioning correctly.

The `/System` folder is the backbone of macOS, housing essential components like the operating system kernel, system frameworks, extensions, drivers, and default applications. These elements ensure macOS runs smoothly and securely by managing system resources, facilitating hardware communication, and providing core functionalities. The folder's read-only nature protects these critical files from unintended changes, maintaining system stability and performance. This organization allows macOS to efficiently manage resources, support hardware interactions, and deliver a consistent user experience, making the `/System` folder a fundamental part of the macOS architecture.



To navigate to any of the macOS system folders, you can open the Finder, then select Go, then Go To Folder, and then type the folder name.

Best Practices

Regardless of the operating system, there are a number of best practices that an administrator should always follow. Depending on the operating system (and distribution, version, edition, and so on), it may be possible to perform operations with the utilities provided, or third-party utilities may be needed.



For Exam Objective 1.8 study, be sure you are familiar with best practices, backups, antivirus, updates/patches, and Rapid Security Response (RSR).

Scheduled Backups

Backups are duplicate copies of critical information, ideally stored in a location other than the one where the information is currently stored. Backups include both paper and computer records. Computer records are usually backed up using a backup program, backup systems, and backup procedures.

The primary starting point for disaster recovery involves keeping current backup copies of important data files, databases, applications, and paper records available for use. Your organization must develop a solid set of procedures to manage this process and to ensure that all critical information is protected.

Computer files and applications should be backed up on a regular basis. Here are some examples of critical files that should be backed up:

- Applications
- Appointment files
- Audit files
- Customer lists
- Database files
- Email correspondence
- Financial data
- Operating systems
- Prospect lists
- Transaction files
- User files
- User information
- Utilities

This list isn't all-inclusive, but it provides a place for you to start.

In most environments, the volume of information that needs to be stored is growing at a tremendous pace. Simply tracking this massive growth can create significant problems.



An unscrupulous attacker can glean as much critical information from copies as they can from the original files. Make sure that your storage facilities are secure, and it is a good idea to add physical security to the backup media as well.

You might need to restore information from backup copies for any number of reasons. Some of the more common reasons are as follows:

- Accidental deletion
- Application errors
- Natural disasters
- Physical attacks
- Server failure
- Virus/malware infection
- Workstation failure

Types of Storage Mechanisms

The information that you back up must be immediately available for use when needed. If a user loses a critical file, they won't want to wait several days while data files are sent from a remote storage facility. Several types of storage mechanisms are available for data storage:

Working Copies Working copy backups, sometimes referred to as *shadow copies*, are partial or full backups that are kept at the computer center for immediate recovery purposes. Working copies are frequently the most recent backups that have been made.

Typically, working copies are intended for immediate use. They are usually updated on a frequent basis.



Working copies aren't usually intended to serve as long-term copies. In a busy environment, they may be created every few hours.

Many filesystems used on servers include journaling. A *journaled file system (JFS)* includes a log file of all changes and transactions that have occurred within a set period of time (such as the last few hours). If a crash occurs, the operating system can check the log files to see which transactions have been committed and which transactions have not.

This technology works well, allowing unsaved data to be written after the recovery, and the system is usually successfully restored to its pre-crash condition.

On-Site Storage *On-site storage* usually refers to a location on the site of the computer center that is used to store information locally. On-site storage containers are available that allow computer cartridges, tapes, and other backup media to be stored in a reasonably protected environment in the building.



As time goes on, tape as a medium for backups is losing its popularity to other technologies, such as cloud storage.

On-site storage containers are designed and rated for fire, moisture, and pressure resistance. These containers aren't fireproof in most situations, but they are *fire-rated*. A *fire-proof container* should be guaranteed to withstand damage regardless of the type of fire or temperature, whereas fire ratings specify that a container can protect the contents for a specific amount of time in a given situation.

If you choose to depend entirely on on-site storage, make sure that the containers you acquire can withstand the worst-case environmental catastrophes that could happen at your location. Make sure, as well, that they are in locations where you can easily find them after the disaster and access them (near exterior walls, on the ground floor, and so forth).



General-purpose storage safes aren't usually suitable for storing electronic media. The fire ratings used for safes generally refer to paper contents. Because paper does not catch fire until 451° Fahrenheit, electronic media are typically ruined well before paper documents are destroyed in a fire.

Off-Site Storage *Off-site storage* refers to a location away from the computer center where paper copies and backup media are kept in the event the primary site suffers a catastrophe. Off-site storage can involve something as simple as keeping a copy of backup media at a *remote office*, or it can be as complicated as a nuclear-hardened high-security storage facility. The storage facility should be bonded, insured, and inspected on a regular basis to ensure that all storage procedures are being followed.

Your determination of which storage mechanism to use should be based on the needs of your organization, the availability of storage facilities, and your budget. Most off-site storage facilities charge based on the amount of space required and the frequency of access needed to the stored information.



Although it is easy to see the need for security at any location where your files are stored, don't overlook the need for security during transportation as well.

Scheduled Disk Maintenance

When files are written to a hard drive, they're not always written contiguously or with all the data located in a single location. When discussing Windows, we talked about Disk Defragmenter, which has existed in almost all versions of Windows, and its ability to take file data that has become spread out over the disk and put it all in the same location, a process known as *defragmenting*. This process decreases the time it takes to retrieve files.


As opposed to FAT- and NTFS-based filesystems, the filesystems used on macOS and Linux rarely, if ever, need to be defragmented. The ext3 and ext4 filesystems are common to Linux, and Apple File System (APFS) is common to macOS. They all have on-the-fly defragmentation methods and implement file allocation strategies differently from their traditional Windows counterparts.

System Updates/App Store

It is important to keep the operating system current and updated. Like Windows, many other operating systems include the ability to update automatically, and almost all can look for updates and tell you when they are available. In the Apple world, the *App Store* represents a location where you can also find updates.

For example, Figure 4.21 shows that a new version of iOS is available on an iPhone. By clicking Learn More, you display the reasons the new version has been released, allowing you to read about the changes and decide whether you want to upgrade. The update notification usually includes a Learn More option and a link to Apple's release page for the iOS update.

FIGURE 4.21 Apple iOS software update



To access the Software Update area on an iPhone or iPad, choose Settings > General > Software Update. However, when an update is available, you will generally see a prompt on your home screen. To access the Software Update section on macOS, choose System Settings from the Apple menu, and then click Software Update. In most cases, unless a production device would be negatively impacted, you should keep systems updated with the latest releases.

Patch Management

As a general rule, updates fix a lot of issues and patches fix a few; multiple patches are rolled into updates. You can't always afford to wait for updates to be released and should install patches—particularly security-related patches—when they are released. Bear in mind that if all the security patches are not installed during the OS installation, attackers can exploit the weaknesses and gain access to information.

A number of tools are available to help with patch management, although the intentions of some are better than others. For example, rather than probe a service remotely and attempt to find a vulnerability, the Nessus vulnerability scanner (www.tenable.com/downloads/nessus) will query the local host to see if a patch for a given vulnerability has been applied. This type of query is far more accurate (and safer) than running a remote check. Since remote checks actually send the exploit in order to check to see if it is applicable, this can sometimes crash a service or process.

Depending on the variant of Linux you are running, *APT (Advanced Package Tool)* can be useful in getting the patches from a repository site and downloading them for installation on Debian and Ubuntu (just to name a few). The most common command used with this tool is `apt update`, which, as the name implies, gets the package for installation, as shown in Figure 4.22. The *Dandified YUM (DNF)* has replaced the *Yellowdog Updater, Modified (YUM)* tool used with *Red Hat Package Manager (RPM)*-based Linux distributions, such as Fedora and Red Hat, and works in a similar way to APT.

FIGURE 4.22 Ubuntu apt tool

```
login as: user
Authenticating with public key
Passphrase for key:
Welcome to Ubuntu 18.04.6 LTS (GNU/Linux 6.10.2-x86_64-linode165 x86_64)

* Documentation: https://help.ubuntu.com

40 updates can be applied immediately.
40 of these updates are ESM Infra security updates.
To see these additional updates run: apt list --upgradable

Last login: Sat Nov  2 14:02:45 2024 from 72.32.46.40
user@srv:~$ sudo apt update
[sudo] password for user:
Hit:1 http://mirrors.linode.com/ubuntu bionic InRelease
Get:2 http://mirrors.linode.com/ubuntu bionic-updates InRelease
Hit:3 http://mirrors.linode.com/ubuntu bionic-backports InRelease
Ign:4 http://apt-longview.linode.com bionic InRelease
Get:5 http://security.ubuntu.com/ubuntu bionic-security InRelease [102 kB]
Get:6 http://apt-longview.linode.com bionic Release
Hit:7 http://ppa.launchpad.net/certbot/certbot/ubuntu bionic InRelease
Get:8 https://esm.ubuntu.com/apps/ubuntu bionic-apps-security InRelease [7,651 B]
Get:9 https://esm.ubuntu.com/apps/ubuntu bionic-apps-updates InRelease [7,554 B]
Get:10 https://esm.ubuntu.com/infra/ubuntu bionic-infra-security InRelease [7,584 B]
Get:11 https://esm.ubuntu.com/infra/ubuntu bionic-infra-updates InRelease [7,526 B]
Reading package lists... Done
E: The repository 'http://apt-longview.linode.com bionic Release' no longer has a Release file.
N: Updating from such a repository can't be done securely, and is therefore disabled by default.
N: See apt-secure(8) manpage for repository creation and user configuration details.
user@srv:~$
```

Driver/Firmware Updates

With any operating system, it is essential to keep the drivers and *firmware* updated. Always remember to back up your configurations (such as for routers) before making any significant changes—in particular, a firmware upgrade—in order to provide a fallback in case something goes awry.

Many network devices contain firmware with which you interact during configuration. For security purposes, you must authenticate in order to make configuration changes and do so initially by using the default account(s). Make sure that the default password is changed after the installation on any network device; otherwise, you are leaving that device open for anyone recognizing the hardware to access it using the known factory password.

Updating firmware for macOS is performed via software updates. During the software update process, the firmware and the software that corresponds to it are updated. Updating firmware for the hardware installed on Linux computers will vary significantly, depending on the type of hardware. Many enterprise Linux vendors, such as Red Hat, include firmware updates in their software updates. As a rule, however, firmware is not part of the Linux software update process.

Antivirus/Antimalware Updates

At one point in time, there were so few viruses outside the Windows world that users not running Windows felt safe without protection on their systems. A significant reason for the low amount of non-Windows *malware* was that the authors of such devious programs were focusing on Windows simply because it had the lion's share of the market; they wanted to inflict as much harm as possible with their code.

As other operating systems have increased in popularity, so too have the number of malware items written for them or that can affect them. Because of this, today it is imperative to have protection on every machine. Additionally, this protection—in the form of definition files—must be kept current and up-to-date. Antivirus and antimalware definitions are released by the hour as new viruses and malware are identified. Most operating systems check daily for updates to definitions. Volume 2 Chapter 5, “Security Concepts,” discusses security and antivirus/antimalware in more detail.

Rapid Security Response (RSR)

Rapid Security Response (RSR) is a feature introduced by Apple to deliver critical security updates to macOS, iOS, and iPadOS devices more swiftly than traditional software updates. These responses are designed to address significant security vulnerabilities that could be exploited by malicious actors, ensuring that devices remain protected against emerging threats.

RSRs can include updates to essential components such as the Safari web browser, the WebKit framework stack, and other critical system libraries. By focusing on these key areas, RSRs can quickly mitigate security issues without waiting for the next full software update cycle. This rapid deployment is crucial for maintaining the security integrity of Apple devices.

One of the advantages of RSRs is their ability to be applied automatically. Devices running the latest versions of the operating system will receive these updates without user intervention, although a restart may be required to complete the installation. This automation ensures that security patches are promptly applied, reducing the window of vulnerability.

RSRs are also designed to be lightweight and efficient. They require less battery power and system resources compared to full software updates, making them less disruptive to the user experience. Additionally, RSRs can be removed if necessary, allowing users to revert to the previous state if any issues arise from the update.

To verify RSRs are automatically installed, follow the following directions for your device:

iOS

1. Navigate to Settings.
2. Select General.
3. Tap Software Update.
4. Choose Automatic Updates and ensure that Security Responses & System Files is selected.

macOS

1. Open the Apple menu.
2. Select System Settings.
3. Click General in the sidebar.
4. Select Software Update on the right.
5. Click the Show Detail button next to Automatic Updates and ensure that Install Security Responses And System Files is selected.

Overall, Rapid Security Response enhances the security posture of Apple devices by providing timely and efficient updates to critical system components, ensuring that users are protected against the latest threats with minimal disruption to their daily activities.

Linux Administration

The best way to approach the following commands is to think about Microsoft Windows. That operating system offers a plethora of utilities for configuring the workstation, and just in case they don't work, or you want to go about it the hard way, you can use command-line utilities to accomplish similar tasks. The odds are good that you spend most of your time walking through graphical dialog boxes, but you're familiar enough with the command-line utilities that you can use them when you need to.

Linux is the same way. There is an overabundance of graphical utilities that can be used to configure the system, and they differ based on the distribution and the graphical interface

being used. In addition to these, command-line utilities are available in every distribution that can be used to get the job done. Those command-line utilities are what we'll focus on here.

There is only one vendor for Windows (Microsoft), but there are many vendors for Linux (Red Hat, SuSE, Ubuntu—to name just three). Also, a new version of Windows is released only every few years (Windows 7, Windows 8/8.1, Windows 10, Windows 11), but with Linux—especially because there are so many vendors—there are a lot of versions. With Ubuntu, for example, the goal is to release a new version every six months.

With all the different distributions and versions, getting to the place where you can run command-line utilities can differ a bit. In almost every implementation of Linux, you can boot into a command-line mode, and the commands entered there can then be run. Better than that, though, the easiest way to get to the command line is to open a terminal (also called a console) window. This allows you to interact with the shell, where you can type commands to your heart's content. The default shell in many Linux distributions is Bash. When you open a terminal window or log in at a text console, the Bash shell is what prompts you for commands. When you type a command, the shell executes your command.



Exam Objective 1.9, “Identify common features and tools of the Linux client/desktop operating system,” requires you to know file management, filesystem management, administrative, package management, network, informational, and text editors used with Linux. In addition, you should know the various common configuration files, and OS components used with Linux.

Understanding the Syntax of Shell Commands

Because a shell interprets what you type, knowing how the shell processes the text you enter is important. All shell commands have the following general format (some commands have no options):

```
command [option1] [option2] ... [optionN]
```

On a command line, you enter a command followed by zero or more *options* (also called *arguments*). The shell uses a blank space or a tab to distinguish between the command and options. This means that you must use a space or a tab to separate the command from the options and the options from one another. If an option contains spaces, you put that option inside quotation marks. For example, to search for a name in the password file, enter the following `grep` command (`grep` is used for searching for text in files):

```
[user@localhost ~]$ grep "Jon B" /etc/passwd
```

When `grep` prints the line with the name, it looks like this:

```
filea:x:1000:100:Jon B:/home/testuser:/bin/bash
[user@localhost ~]$
```

If you create a user account with your username, type the `grep` command with your username as an argument to look for that username in the `/etc/passwd` file. In the output from the `grep` command, you can see the name of the shell (`/bin/bash`) following the last colon (`:`). Because the Bash shell is an executable file, it resides in the `/bin` directory; you must provide the full path to it.

The number of command-line options and their formats depend on the actual command. Typically, these options look like `-X`, where `X` is a single character. For example, you can use the `-l` option with the `ls` command. The command lists the contents of a directory, and the option provides additional details. Here is a result of typing `ls -l` in a user's home directory:

```
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x 2 user users 48  2018-09-08 21:11 bin
drwx-      2 user users 320 2018-09-08 21:16 Desktop
drwx-      2 user users 80   2018-09-08 21:11 Documents
drwxr-xr-x 2 user users 80   2018-09-08 21:11 public_html
drwxr-xr-x 2 user users 464 2018-09-17 18:21 sdump
[user@localhost ~]$
```



The `/etc/passwd` file contains a list of the usernames and their respective attributes, such as User ID (UID), Group ID (GID), User ID info, home directory, and shell. Interestingly enough, it doesn't normally store the user's password; another file, called the shadow file (`/etc/shadow`), now contains a hashed version of the user's password.

If a command is too long to fit on a single line, you can press the backslash key (`\`) followed by Enter. Then continue typing the command on the next line. For example, type the following command (press Enter after each line):

```
[user@localhost ~]$ cat \
> /etc/passwd
root:x:0:0:Super User:/root:/bin/bash
bin:x:1:1:bin:/bin:/usr/sbin/nologin
daemon:x:2:2:daemon:/sbin:/usr/sbin/nologin
adm:x:3:4:adm:/var/adm:/usr/sbin/nologin
[Output Cut]

[user@localhost ~]$
```

The `cat` command then displays the contents of the `/etc/passwd` file.

You can concatenate (that is, string together) several shorter commands on a single line by separating the commands with semicolons (`;`). For example, the following command

changes the *current directory* to your home directory, lists the contents of that directory, and then shows the name of that directory:

```
[user@localhost ~]$ cd /dev; ls -l; pwd
total 0
crw-r--r--. 1 root root      10, 235 Jan  1 12:28 autofs
drwxr-xr-x. 2 root root      180 Jan  1 12:28 block
drwxr-xr-x. 2 root root       60 Jan  1 12:28 bsg
crw-rw----. 1 root disk     10, 234 Jan  1 12:28 btrfs-control
drwxr-xr-x. 3 root root       60 Jan  1 12:28 bus
[Output Cut]
```

```
[user@localhost dev]$
```

You can combine simple shell commands to create a more sophisticated command. For example, suppose you want to find out whether a device file named `nvme0` resides in your system's `/dev` directory because some documentation says that you need that device file for your hard drive. You can use the `ls /dev` command to get a directory listing of the `/dev` directory, and then browse through it to see whether that listing contains `nvme0`.

Unfortunately, the `/dev` directory has a great many entries, so you may find it hard to find any item that has `nvme0` in its name. You can, however, combine the `ls` command with `grep` and come up with a command line that does exactly what you want. Here's that command line:

```
[user@localhost ~]$ ls /dev | grep nvme0
nvme0
nvme0n1
nvme0n1p1
nvme0n1p2
nvme0n1p3
[user@localhost ~]$
```

The shell sends the output of the `ls` command (the directory listing) to the `grep` command, which searches for the string `nvme0`. The vertical bar (`|`) is known as a *pipe* because it acts as a conduit (think of a water pipe) between the two programs—the output of the first command is fed into the input of the second command.

Discovering and Using Linux Commands

Literally hundreds, if not thousands, of Linux commands exist within the shell and the system directories. Fortunately, CompTIA asks that you know a much smaller number than that. Table 4.1 lists common Linux commands by category.

TABLE 4.1 Essential Linux commands

Command name	Action
Managing files and directories	
cd	Changes the current directory.
chmod	Changes file permissions.
chown	Changes the file owner and group.
cp	Copies files.
ls	Displays the contents of a directory.
md	Creates a directory.
mv	Renames a file and moves the file from one directory to another.
rm	Deletes files.
pwd	Displays the current directory.
find	Searches for text in a file hierarchy.
grep	Searches for regular expressions in a text file.
Filesystem management	
fsck	Fixes the filesystem for issues or corruption.
mount	Utility used to mount folders and drives to the filesystem.
Administrative	
passwd	Changes the password.
su	Starts a new shell as another user. (The other user is assumed to be root when the command is invoked without any argument.)
sudo	Runs a command as another user (usually the root user).
shutdown	Shuts down Linux.

Command name	Action
Package management	
apt	Downloads files from a repository site, used with Debian-based Linux.
dnf	Downloads files from a repository site, used with Red Hat-based Linux.
Network	
dig	DNS query utility.
ip	Allows you to display and configure information related to a network interface card (NIC).
curl	Allows transferring data between servers using various protocols like HTTP and FTP.
ping	Used to test the reachability of a network host and measure the round-trip time for a packet.
traceroute	Used to map the path data packets take from your computer to a destination server.
Text editors	
vi	Starts the visual file editor, which can be used to edit files.
nano/pico	Starts the visual file editor, which is simple to use for editing files.
Informational	
man	System help for executable files.
df	Displays the total disk free (disk free space) for a directory.
dd	Copies blocks of data from one file to another (used to copy data from devices).
du	Allows users to analyze and report on disk usage within directories and files.
ps	Displays a list of currently running processes.
kill	Terminates a process.

TABLE 4.1 Essential Linux commands *(continued)*

Command name	Action
cat	Displays the contents of a file.
top	Displays running processes, similar to Windows Task Manager.

User Accounts

Linux is a powerful and versatile operating system widely used in various environments, from personal computers to enterprise servers. One of the fundamental principles of Linux system administration is the creation of non-root users during the initial setup. The root user, also known as the superuser, has unrestricted access to all system commands and files, making it a highly privileged account. While this level of access is necessary for certain administrative tasks, it also poses significant risks if misused or compromised.

Creating non-root users allows for the delegation of tasks and limits the potential damage that can be caused by accidental or malicious actions. Non-root users have restricted permissions, a fact that helps prevent unauthorized changes to critical system files and settings. This separation of privileges is a key aspect of the principle of least privilege concept, which aims to minimize the risk of security breaches and system failures.

For example, consider a scenario where a user accidentally runs a command that deletes important system files. If the user is operating as root, the command will execute with full privileges, potentially rendering the system inoperable. However, if the user is a non-root user, the command will fail due to insufficient permissions, protecting the system from unintended harm. This example illustrates why it is strongly discouraged to run as root for everyday tasks and highlights the importance of creating and using non-root user accounts in Linux.

Becoming Root (Superuser)

When you want to do anything that requires a high privilege level (e.g., administering your system), you have to become *root*. Normally, you log in as a regular user with your everyday username. When you need the privileges of the *superuser*, though, use the following command to become root:

```
[user@localhost ~]$ whoami
user
[user@localhost ~]$ su -
Password:
[root@localhost ~]# whoami
root
[root@localhost ~]#
```

The `su` command followed by a space and the minus sign (or hyphen) provides an environment similar to what the user would expect by applying the user's environment variables and initial login scripts. Once executed, the shell then prompts you for the root password. Type the password and press Enter. Your command prompt will change to a `#` symbol and your security context will change to the root user. The `whoami` command is a life saver, it helps you keep your sanity since the prompts look so familiar.



When the `su` command is entered without a username as an option, you will be prompted for the password of the root account, and a Bash shell will open on behalf of the root account. However, you can use the `su` command followed by any user account name, and as long as you supply the correct password, you will be logged in as that user.

After you've finished with whatever you want to do as root (and you have the privilege to do anything as root), type `exit` to return to your normal username.

Instead of becoming root by using the `su -` command, you can type **`sudo`** followed by the command that you want to run as root. In some distributions, such as Ubuntu, you must use the `sudo` command because you don't get to set up a root user when you install the operating system. If you're listed as an authorized user in the `/etc/sudoers` file, `sudo` executes the command as if you were logged in as root. Type **`man sudoers`** to read more about the `/etc/sudoers` file.



As you learn many of the commands in this chapter, you'll realize that it is similar to learning a new language. Therefore, practice makes perfect, and we urge you to practice these commands. The `man` command can help with nuances of Linux. The command provides a manual pages for a command or component installed on the Linux operating system. As new commands are installed the man pages are added.

Creating Accounts

Creating a new user account in Linux is a straightforward process that involves using the `useradd` command followed by the desired username. For example, to create a user named *CompTIA*, you would execute the following command:

```
user@server:~$ sudo useradd CompTIA
[sudo] password for user:
user@server:~$
```

After creating the account, it's essential to set a password for the new user using the `passwd` command:

```
user@server:~$ sudo passwd CompTIA
[sudo] password for user:
New password:
```

```

Retype new password:
passwd: password updated successfully
user@server:~$

```

The `sudo` command will prompt you to enter your password first, then enter and confirm the new password for the user `CompTIA`. Additionally, you can customize the user's home directory, shell, and other settings by using various options with the `useradd` command. Creating user accounts allows for better management of system resources and enhances security by assigning specific permissions and roles to different users.

Users, Groups, and Passwords

The `/etc/passwd` file is a crucial component of user account management in Linux. It contains essential information about each user, including the username, user ID (UID), group ID (GID), home directory, and default shell. Each line in the file represents a single user account, with fields separated by colons. Historically, this file also stored user passwords in an encrypted format. However, to enhance security, modern systems now use the `/etc/shadow` file for storing password information. The `/etc/passwd` file remains publicly readable, allowing various system utilities to access user information while keeping sensitive data secure. The following is a typical `/etc/passwd` file with the addition of our initial user account and the `CompTIA` user.

```

[user@localhost ~]$ cat /etc/passwd
root:x:0:0:root:/root:/bin/bash
bin:x:1:1:bin:/bin:/usr/sbin/nologin
daemon:x:2:2:daemon:/sbin:/usr/sbin/nologin
adm:x:3:4:adm:/var/adm:/usr/sbin/nologin
lp:x:4:7:lp:/var/spool/lpd:/usr/sbin/nologin
[Output Cut]

user:x:1000:1000:user:/home/user:/bin/bash
CompTIA:x:1001:1001::/home/CompTIA:/bin/bash
[user@localhost ~]$

```

In addition to the user account created in the `/etc/passwd` file, a group for the user is also created in the `/etc/group` file. You can add users to groups by using the command `usermod -aG {group name} {user}`. As an example, in the following we will add the `CompTIA` user to the `root` group:

```

[user@localhost ~]$ sudo usermod -aG root CompTIA
[sudo] password for user:
[user@localhost ~]$ cat /etc/group
root:x:0:CompTIA
bin:x:1:
daemon:x:2:
sys:x:3:

```



```
adm:x:4:
tty:x:5:
disk:x:6:
lp:x:7:
[Output Cut]
```

```
[user@localhost ~]$
```

The `/etc/shadow` file is specifically designed to securely store user password information. Unlike the `/etc/passwd` file, it is only accessible by the root user and contains hashed passwords along with additional fields for password aging and expiration policies. By separating password data from the more publicly accessible `/etc/passwd` file, Linux systems significantly improve the security of user credentials. The `/etc/shadow` file ensures that sensitive password information is protected from unauthorized access, reducing the risk of security breaches and enhancing overall system security. Proper management of this file is essential for maintaining a secure Linux environment.

```
[user@localhost ~]$ cat /etc/shadow
cat: /etc/shadow: Permission denied
[user@localhost ~]$ sudo cat /etc/shadow
[sudo] password for user:
root:$y$j9T$JT0Rk.nzkReN8x6d/40BfUGj$RR9uljt9q0mWJP9hbtFwtEvaEynG3CBfoRQEhuq
qse1::0:99999:7:::
bin:*:19925:0:99999:7:::
daemon:*:19925:0:99999:7:::
adm:*:19925:0:99999:7:::
lp:*:19925:0:99999:7:::
[Output Cut]

user:$y$j9T$JAnsgsv9zypBiyL/bKgX0eCP$JW97Kaq0ym07nTMGdKGitK10JzhQYSBqk7
.UF3i0Ef9::0:99999:7:::
CompTIA:$y$j9T$e5R5vtTEfA25/HIQboCB81$TnGiEWcdTZsHpHsHhoSe3J6zi7055e5l8
.QXxcqPWC9:20091:0:99999:7:::
[user@localhost ~]$
```

Maintaining the Operating System

The Linux operating system is a secure and functional operating system. However, it's inevitable that at some point you will need to obtain updates for security or to extend the functionality by adding packages. Thankfully, every Linux distribution has its own unique repository of security patches and additional packages. When choosing a Linux distribution, you should spend time exploring its benefits. For example, Debian packages are more stable, but the disadvantage is that the packages are older than the current version of packages.

Mint Linux is tailored to laptops and desktops, but it is not normally found on servers. These are just a few examples—the lesson is to do your homework before you decide on a Linux distribution.

Once you have chosen and installed a Linux distribution, you'll want to update the repositories and then upgrade the distribution to get the latest security patches. A *repository* is a group of packages that are available for download. The repository contains metadata, such as versions, dates, descriptions, and dependencies, about these packages. When you update the repositories, you are downloading the metadata so that you can search and begin the upgrade process.

Depending on the version, you will have one of two or three tools: the Advanced Package Tool (APT), Dandified YUM (DNF), or Yellowdog Updater, Modified (YUM). Linux distributions, such as Ubuntu, Debian, and Mint, will use the APT package management tool. In the operating system, you can update the repositories by using the `apt` or the `apt-get` command. In addition to upgrading the operating system, you can install and manage packages. In order to update the operating system, you will need to update the repositories first and then you can upgrade the operating system binaries as shown in the following:

```
user@server:~$ sudo apt update
[sudo] password for user:
Hit:1 http://archive.ubuntu.com/ubuntu bionic InRelease
Hit:2 http://archive.ubuntu.com/ubuntu bionic-updates InRelease
Hit:3 http://archive.ubuntu.com/ubuntu bionic-backports InRelease
Hit:4 http://archive.ubuntu.com/ubuntu bionic-security InRelease
Reading package lists... Done
user@server:~$ sudo apt upgrade
Reading package lists... Done
Building dependency tree
Reading state information... Done
Calculating upgrade... Done
The following packages have been kept back:
  base-files netplan.io sosreport ubuntu-advantage-tools ubuntu-server
The following packages will be upgraded:
  accountsservice apport apt apt-utils bash bcache-tools bind9-host bsdtls
[ output cut]
  unattended-upgrades update-manager-core update-notifier-common ureadahead
  util-linux uuid-runtime vim vim-common vim-runtime vim-tiny wget xkb-data
241 upgraded, 0 newly installed, 0 to remove and 5 not upgraded.
Need to get 111 MB of archives.
After this operation, 43.2 MB of additional disk space will be used.
Do you want to continue? [Y/n]
```

In the following example, you can see how a package such as `iftop` is installed using the `apt` command:

```
user@server:~$ sudo apt install iftop
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following NEW packages will be installed:
  iftop
0 upgraded, 1 newly installed, 0 to remove and 246 not upgraded.
Need to get 36.0 kB of archives.
After this operation, 91.1 kB of additional disk space will be used.
Get:1 http://archive.ubuntu.com/ubuntu bionic/universe amd64 iftop amd64
1.0~pre4-4 [36.0 kB]
Fetched 36.0 kB in 0s (102 kB/s)
Selecting previously unselected package iftop.
(Reading database ... 66991 files and directories currently installed.)
Preparing to unpack .../iftop_1.0~pre4-4_amd64.deb ...
Unpacking iftop (1.0~pre4-4) ...
Setting up iftop (1.0~pre4-4) ...
Processing triggers for man-db (2.8.3-2ubuntu0.1) ...
user@server:~$
```

The DNF and YUM package managers are used to update and install packages for Red Hat-based Linux distributions, such as Red Hat Enterprise Server and Fedora, among others. These tools work just like the APT tool; the first step is to update the repositories, then you can update the binaries. In the following examples, we will demonstrate both utilities:

```
[root@localhost ~]# yum update
CentOS Stream 8 - AppStream      5.6 MB/s | 16 MB 00:02
CentOS Stream 8 - BaseOS        1.6 MB/s |  6 MB 00:02
CentOS Stream 8 - Extras        35 kB/s | 15 kB 00:00
Dependencies resolved.
Nothing to do.
Complete!
[root@localhost ~]# yum upgrade
Last metadata expiration check: 0:01:40 ago on Fri 22 Oct 2021
09:52:10 PM EDT.
Dependencies resolved.
Nothing to do.
Complete!
[root@localhost ~]#
```

The yum command can also be used to install packages. In the following example, we are using yum to install the nano utility:

```
[root@localhost ~]# yum install nano
Last metadata expiration check: 0:45:41 ago on Fri 22 Oct 2021
10:51:16 PM EDT.
Dependencies resolved.
=====
Package           Architecture      Version           Repository        Size
=====
Installing:
 nano             x86_64            2.9.8-1.el8       baseos            581 k

Transaction Summary
=====
Install 1 Package

Total download size: 581 k
Installed size: 2.2 M
Is this ok [y/N]:
```



The yum command has been dropped as an exam objective from the CompTIA exam. However, it is important to know it exists and functions the same as the apt and dnf commands. Therefore, it has been included for completeness. Linux is incredibly stable once it is installed and you could find yourself upgrading a system that only has the yum command available.

Newer Red Hat Linux distributions now come with the DNF package manager. You will love this tool! One simple command both updates the repositories and upgrades them. The dnf command actually looks really similar to yum in its output. You might even have both commands on the Linux distribution. Let's see the typical output of dnf updating and upgrading a Fedora installation in the following. We cut some of the output since the update/upgrade would take several pages. Anywhere there is [Output Cut] signifies we eliminated redundant output for simplicity's sake.

```
[user@localhost ~]$ sudo dnf upgrade
[sudo] password for user:
Updating and loading repositories:
Fedora 41 openh264 (From Cisco) - x86_ 100% | 11.8 KiB/s | 4.8 KiB | 00m00s
Fedora 41 - x86_64 - Updates           100% | 7.9 MiB/s | 7.9 MiB | 00m01s
Fedora 41 - x86_64                     100% | 16.1 MiB/s | 36.4 MiB | 00m02s
Repositories loaded.
```

```

Package                               Arch  Version
Reposit      Size
Upgrading:
  PackageKit                           x86_64 1.2.8-8.fc41
updates    2.9 MiB
    replacing PackageKit                x86_64 1.2.8-7.fc41
anacond    2.9 MiB
  PackageKit-glib                       x86_64 1.2.8-8.fc41
updates   520.3 KiB
    replacing PackageKit-glib           x86_64 1.2.8-7.fc41
[Output Cut]

Installing:
  kernel                               x86_64 6.12.6-200.fc41
updates    0.0  B
Installing dependencies:
  hiredis                              x86_64 1.2.0-3.fc41
fedora   110.1 KiB
  kernel-core                           x86_64 6.12.6-200.fc41
updates    73.6 MiB
  kernel-modules                         x86_64 6.12.6-200.fc41
[Output Cut]

Transaction Summary:
Installing:      9 packages
Upgrading:      190 packages
Replacing:      190 package

Total size of inbound packages is 518 MiB. Need to download 518 MiB.
After this operation, 182 MiB extra will be used (install 871 MiB, remove 688 MiB).
Is this ok [y/N]: y
[ 1/199] kernel-0:6.12.6-200.fc41.x86_ 100% | 499.0 KiB/s | 191.1 KiB | 00m00s
[ 2/199] kernel-core-0:6.12.6-200.fc41 100% | 11.0 MiB/s | 18.4 MiB | 00m02s
[ 3/199] opensc-libs-0:0.26.0-1.fc41.x 100% | 10.7 MiB/s | 918.0 KiB | 00m00s
[Output Cut]

-----
[199/199] Total                               100% | 39.6 MiB/s | 518.0 MiB | 00m13s
Running transaction
[ 1/391] Verify package files                 100% | 153.0  B/s | 199.0  B | 00m01s
[ 2/391] Prepare transaction                 100% | 1.0 KiB/s | 389.0  B | 00m00s
[ 3/391] Upgrading glibc-common-0:2.40 100% | 12.5 MiB/s | 1.0 MiB | 00m00s

```

[Output Cut]

```
[202/391] Erasing abrt-cli-0:175:2.17.6-2.f 100% | 97.7 KiB/s | 100.0 B | 00m00s
[203/391] Erasing abrt-tui-0:2.17.6-2.f 100% | 25.9 KiB/s | 53.0 B | 00m00s
[204/391] Erasing python3-dnf-0:4.21.1- 100% | 67.1 KiB/s | 275.0 B | 00m00s
[Output Cut]
```

Complete!

[user@localhost ~]\$

Let's examine how packages are added with the `dnf` command. You'll find it identical to `apt` and `yum`. In the following output, we will install the `nano` package with the command `sudo dnf install nano`.

[user@localhost ~]\$ **sudo dnf install nano**

Updating and loading repositories:

Repositories loaded.

Package	Arch	Version	Repository	Size
Installing:				
nano	x86_64	8.1-1.fc41	fedora	2.9 MiB

Transaction Summary:

Installing: 1 package

Total size of inbound packages is 729 KiB. Need to download 729 KiB.

After this operation, 3 MiB extra will be used (install 3 MiB, remove 0 B).

Is this ok [y/N]: y

```
[1/1] nano-0:8.1-1.fc41.x86_64 100% | 3.3 MiB/s | 728.9 KiB | 00m00s
```

```
-----
[1/1] Total 100% | 1.8 MiB/s | 728.9 KiB | 00m00s
```

Running transaction

```
[1/3] Verify package files 100% | 200.0 B/s | 1.0 B | 00m00s
```

```
[2/3] Prepare transaction 100% | 15.0 B/s | 1.0 B | 00m00s
```

```
[3/3] Installing nano-0:8.1-1.fc41.x86_ 100% | 16.4 MiB/s | 2.9 MiB | 00m00s
```

Complete!

[user@localhost ~]\$

Now that you know how to update the operating system and install packages, it's inevitable that you'll run out of space on the operating system. Fortunately, you can monitor and quickly find out how much space is free on the disk by using the `df` command, otherwise known as the disk free command. By using the `df` command you can quickly see the percentage of free space, and if you supply the `-h` argument you'll get results in human-readable formats of bytes, as shown here:

user@server:~\$ **df -h**

Filesystem	Size	Used	Avail	Use%	Mounted on
------------	------	------	-------	------	------------

```
udev                1.9G      0  1.9G    0% /dev
tmpfs               393M    1.5M  391M    1% /run
/dev/sda2           20G    6.0G   13G   33% /
tmpfs               2.0G      0  2.0G    0% /dev/shm
tmpfs               5.0M      0  5.0M    0% /run/lock
tmpfs               2.0G      0  2.0G    0% /sys/fs/cgroup
/dev/loop0          91M    91M      0 100% /snap/core/6350
/dev/loop1          100M   100M      0 100% /snap/core/11993
tmpfs               393M      0  393M    0% /run/user/1000
user@server:~$
```

Filesystem Management

The filesystem is a critical component of any Linux operating system. This section will cover some key aspects of filesystem management in Linux that CompTIA has identified as exam objectives for the 220-1202 exam.

mount Command

The `mount` command in Linux is a fundamental utility used to attach filesystems to the directory tree, making them accessible to the system and users. This command allows you to mount various types of filesystems, including local disk partitions, network shares, and removable media like USB drives. By specifying the device and the mount point, you can integrate additional storage into your system seamlessly. For example, to mount a USB drive located at `/dev/sdb1` to the `/mnt/usb` directory, you can use the following command:

```
[user@localhost ~]$ sudo mount -t vfat /dev/sdb1 /mnt/usb
[sudo] password for user:
[user@localhost folder]$ ls -l /mnt/usb
total 0
-rw-rw-rw-. 1 user user 44 Jan  1 13:07 pinball.exe
-rw-rw-rw-. 1 user user 34 Jan  1 12:02 file1.txt
-rw-rw-rw-. 1 user user 34 Jan  1 14:15 file2.txt
-rw-rw-rw-. 1 user user 36 Jan  1 15:05 file3.txt
[user@localhost folder]$
```

The `mount` command by itself will display all the current mount points on the operating system. This list is pretty extensive on a modern Linux operating system. So, we will use the `grep` command to only return the result for our mount of `sdb1`, as shown in the following:

```
[user@localhost ~]$ mount | grep sdb1
/dev/sdb1 on /mnt/usb type vfat (rw,nosuid,nodev,relatime,user_id=0,group_id=0,allow_other,blksize=4096,x-gvfs-show)
[user@localhost ~]$
```

The `mount` command can also mount an ISO image to the filesystem with something called a *loop device*. A loop device is nothing more than a device driver that acts as a block storage device. `mount` typically mounts devices and not other filesystems. This command would mount an ISO file to the `/mnt/iso` folder:

```
[user@localhost folder]$ sudo mount -o loop ~/file.iso /mnt/iso
[sudo] password for user:
[user@localhost folder]$
```

The `/etc/fstab` configuration file

The `/etc/fstab` file, short for *filesystems table*, complements the `mount` command by automating the mounting process. This critical configuration file defines how disk partitions, storage devices, and other filesystems are mounted and integrated into the system's directory structure at boot time. Each entry in `/etc/fstab` specifies the device, mount point, filesystem type, and various mount options, ensuring consistent and reliable access to storage resources across reboots. Here's an example of `/etc/fstab`:

```
[user@localhost folder]$ sudo cat /etc/fstab
[sudo] password for user:
# <file system> <mount point> <type> <options> <dump> <pass>
/dev/disk/by-id/dm-uuid-LVM-0j77W3g2yUV2jWGh7e8wL7rkCD3wItJYdw4mvlia987HPTxAzBeKjavYcK7fZKAz / ext4 defaults 0 0
/dev/disk/by-uuid/a4bdbde5-70cf-4adf-b94f-e5173310e076 /boot ext4 defaults 0 0
/swap.img none swap sw 0 0
/dev/disk/by-id/usb-G-DRIVE_USB-0:0-part2 /media/Storage auto rw,nosuid,nodev,nofail,x-gvfs-show 0 0
[user@localhost folder]$
```

In the `/etc/fstab` file, each line mounts a different portion of the filesystem. By looking at the second argument of the first line, we can see this is the mount point for `/`, also known as the root of the filesystem. The first argument explains what to mount to the root—in this case, it is a universally unique ID (UUID) used by the Logical Volume Manager (LVM) service. Looking at the third argument, we see it has the `ext4` filesystem. The fourth argument contains the options to be used—in this case, we are using the default. The fifth argument specifies whether the mount point should be backed up, otherwise known as *dumped*. The sixth argument tells the OS whether during bootup the filesystem needs to have filesystem consistency checked.

Creating entries in `/etc/fstab` is a bit like baking at high altitude; it takes some time to get use to the variables. A trick is to edit your line in `/etc/fstab`, then use the command `mount -av` to remount all mount points in the `/etc/fstab`.

Now let's examine the second line in the example `/etc/fstab` file. The device UUID of `/dev/disk/by-uuid/a4bdbde5-70cf-4adf-b94f-e5173310e076` is mounted to the `/boot` folder. This is significant because the `/boot` mount point contains both the kernel and the bootloader configuration. The kernel will be a variation of `vmLinux-*` files, and the bootloader configuration in the following example is the GRand Unified Bootloader (GRUB)

bootloader v2. However, some Linux distributions might also use the Linux LOader (LILO) bootloader.

```
[user@localhost ~]$ ls /boot/
config-6.11.4-301.fc41.x86_64
config-6.12.6-200.fc41.x86_64
efi
grub2
initramfs-0-rescue-12d7e45e4aa444d88c33a0bc2106e0d2.img
initramfs-6.11.4-301.fc41.x86_64.img
initramfs-6.12.6-200.fc41.x86_64.img
loader
symvers-6.11.4-301.fc41.x86_64.xz
symvers-6.12.6-200.fc41.x86_64.xz
System.map-6.11.4-301.fc41.x86_64
System.map-6.12.6-200.fc41.x86_64
vmlinuz-0-rescue-12d7e45e4aa444d88c33a0bc2106e0d2
vmlinuz-6.11.4-301.fc41.x86_64
vmlinuz-6.12.6-200.fc41.x86_64
[user@localhost ~]$
```

The *fsck* Command

The *fsck* (filesystem consistency check) command in Linux is a crucial tool for maintaining the integrity of filesystems. It is used to check and repair inconsistencies and errors in filesystems, ensuring that data remains accessible and the system operates smoothly. The *fsck* command can be run manually or automatically during system boot, depending on the configuration. Here's how you can use *fsck* to check and repair a filesystem:

1. First, identify the device name of the filesystem you want to check. You can use the `lsblk` or `df -h` command. Let's use the example of `/dev/sda1`.

```
[user@localhost ~]$ sudo lsblk
[sudo] password for user:
NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sr0                  11:0    1   2.6G  0 rom
zram0                252:0    0   3.8G  0 disk [SWAP]
nvme0n1              259:0    0   40G   0 disk
|-nvme0n1p1          259:1    0    1M   0 part
|-nvme0n1p2          259:2    0    1G   0 part /boot
|-nvme0n1p3          259:3    0   39G   0 part
| -fedora-root        253:0    0    15G   0 lvm  /
sdb                  260:0    0    8G   0 disk
|-sdb1                260:1    0   39G   0 part /mnt/usb
[user@localhost ~]$
```

2. Before running the `fsck` command, ensure the filesystem is unmounted to prevent data corruption:

```
[user@localhost ~]$ sudo umount /dev/sda1
[sudo] password for user:
[user@localhost ~]$
```

3. Use the `fsck` command to check and repair the filesystem. You can add the `-y` option to automatically answer yes to all prompts:

```
[user@localhost ~]$ sudo fsck -y /dev/sda1
[sudo] password for user:
fsck from util-linux 2.40.2
e2fsck 1.46.5 (30-Dec-0040)
Files: clean, 1129/1200 files, 80340/162345 blocks
[user@localhost ~]$
```

4. The `fsck` command will display the results of the check and any repairs it performed. Review these results. The rule of thumb is that no errors is good news!
5. After the check and repair, remount the filesystem:

```
[user@localhost ~]$ sudo mount /dev/sdb1 /mnt/usb
[sudo] password for user:
[user@localhost ~]$
```

The Linux Boot Process

The Linux boot process, much like its Windows counterpart, is a sequence of events that transforms a powered-off machine into a fully operational system. While Windows relies on the BIOS/UEFI to load the Windows Boot Manager and subsequently the Windows OS loader, Linux employs a combination of EFI (Extensible Firmware Interface), GRUB (GRand Unified Bootloader), and the Linux kernel. Understanding these processes is essential for anyone looking to delve deeper into system administration or operating system internals. The following is a step-by-step breakdown of how these components work together to boot a Linux system:

1. When the computer is powered on, the EFI firmware initializes the hardware and performs a series of checks to ensure everything is functioning correctly. It then looks for the EFI bootloader on the EFI System Partition (ESP).
2. The EFI firmware loads the GRUB bootloader from the ESP. GRUB provides a user interface that allows you to select which operating system or kernel to boot if multiple options are available. It also loads its configuration file, typically `grub.cfg`, which contains the boot menu and options.
3. Once a selection is made, GRUB loads the chosen Linux kernel into memory. It also loads the initial RAM disk (`initramfs` or `initrd`), which contains essential drivers and tools needed to mount the root filesystem.

4. The kernel decompresses itself and initializes the hardware, including the CPU, memory, and peripheral devices. It then mounts the `initramfs` as a temporary root filesystem.
5. The kernel starts the `init` system, such as `systemd` or the traditional `init`. This system is responsible for starting and managing system services and user processes. It brings the system to a specific runlevel or target, defining the state of the system (e.g., multiuser mode, graphical mode).
6. Finally, the system reaches a state where it can present a login prompt or graphical login screen to the user. After login, the user session is started, and the system is ready for use.

While the Linux boot process involves unique components like EFI, GRUB, and the Linux kernel, it shares fundamental similarities with the Windows boot process. Both systems rely on firmware initialization, a bootloader, and a kernel to transition from power-on to a fully functional operating environment. Being familiar with these parallels helps you better understand the Linux operating system by unshrouding the mystery under the hood (so to speak).

Managing Processes

Every time the shell executes a command that you type, it starts a process. The shell itself is a process, as are any scripts or programs that the shell runs. The `ps` command will show a snapshot of the current processes running as the currently logged-in user. Use the `ps ax` command to see a list of processes for the entire operating system. When you type `ps ax`, Bash shows you the current set of processes. Here are a few lines of output from the command `ps ax --cols 80` (the `--cols 80` option is used to ensure that you see each command in its entirety):

```
[user@localhost ~]$ ps ax --cols 80
PID TTY      STAT   TIME COMMAND
   1 ?        Ss      0:03 /usr/lib/systemd/systemd --system
   2 ?        S        0:00 [kthreadd]
   3 ?        S        0:00 [pool_workqueue_release]
   4 ?        I<       0:00 [kworker/R-rcu_gp]
   5 ?        I<       0:00 [kworker/R-sync_wq]
   6 ?        I<       0:00 [kworker/R-slab_flushwq]
   7 ?        I<       0:00 [kworker/R-netns]
  10 ?        I<       0:00 [kworker/0:0H-kblockd]
[Output Cut]
12303 ?        S        0:00 systemd-nsresourcework: waiting...
12304 ?        S        0:00 systemd-nsresourcework: waiting...
12319 ?        S        0:00 systemd-nsresourcework: waiting...
12321 ?        I        0:00 [kworker/0:2-events]
12341 ?        I        0:00 [kworker/1:0-events]
12342 ?        S        0:00 systemd-userwork: waiting...
12345 pts/0    R+       0:00 ps -ax --cols 80
[user@localhost ~]$
```

In this listing, the first column has the heading PID, and it shows a number for each process. PID stands for *process ID (identification)*, which is a sequential number assigned by the Linux kernel. If you look through the output of the `ps ax` command, you'll see that the `systemd` command is the first process and has a PID of 1. That's why `systemd` is referred to as the *mother of all processes*.

The `COMMAND` column shows the command that created each process, and the `TIME` column shows the cumulative CPU time used by the process. The `STAT` column shows the state of a process: S means that the process is sleeping, and R means that it's running. The symbols following the status letter have further meanings; for example, < indicates a high-priority process, and + means that the process is running in the foreground. The `TTY` column shows the terminal, if any, associated with the process.

The process ID, or process number, is useful when you have to stop an errant process forcibly. Look at the output of the `ps ax` command and note the PID of the offending process. Then use the `kill` command with that process number to stop the process. For example, to stop process number 8550, type the following command:

```
[user@localhost ~]$ kill 8550
[user@localhost ~]$
```

The `ps` command will allow you to see processes currently running at the time you initiate the command. However, if you want to see an interactive display of processes similar to the Windows Task Manager, you can use the `top` command, shown in Figure 4.23. The `top` command allows you to sort by columns and scroll through the various processes. You can even kill processes interactively.

As you can see, the `top` command is extremely similar to the Windows Task Manager. Some of the common uses for the `top` command are similar to its Windows counterpart:


Monitoring System Performance The `top` command provides a real-time overview of system performance by displaying key metrics such as CPU and memory usage. This allows users to quickly identify any processes that are consuming excessive resources, which can help in diagnosing performance issues and ensuring that the system runs smoothly. By continuously updating the display, `top` offers an up-to-date snapshot of the system's health.

Process Management With `top`, users can view detailed information about each running process, including the process ID (PID), user, priority, and execution time. This information is crucial for managing processes, especially when troubleshooting issues or optimizing system performance. For instance, if a particular process is causing problems, users can easily identify it and take appropriate action, such as terminating or renicing the process.

Resource Allocation The `top` command helps users understand how system resources are being allocated by displaying resource usage for each process. This includes metrics like CPU and memory consumption, which are essential for effective resource management. By analyzing this data, users can make informed decisions about resource

allocation, such as adjusting the priority of processes or redistributing workloads to ensure optimal performance.

FIGURE 4.23 The top command



```

top - 03:31:26 up 1:52, 1 user, load average: 0.00, 0.00, 0.00
Tasks: 160 total, 1 running, 85 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.0 us, 0.0 sy, 0.0 ni, 100.0 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 4015808 total, 2263564 free, 204664 used, 1547580 buff/cache
KiB Swap: 4015100 total, 4015100 free, 0 used, 3524864 avail Mem
  
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1	root	20	0	159952	9496	6992	S	0.0	0.2	0:04.00	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.01	kthreadd
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	kworker/0:0H
6	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	mm_percpu_wq
7	root	20	0	0	0	0	S	0.0	0.0	0:00.40	ksoftirqd/0
8	root	20	0	0	0	0	I	0.0	0.0	0:00.60	rcu_sched
9	root	20	0	0	0	0	I	0.0	0.0	0:00.00	rcu_bh
10	root	rt	0	0	0	0	S	0.0	0.0	0:00.03	migration/0
11	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/0
12	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/0
13	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/1
14	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	watchdog/1
15	root	rt	0	0	0	0	S	0.0	0.0	0:00.07	migration/1
16	root	20	0	0	0	0	S	0.0	0.0	0:00.43	ksoftirqd/1
18	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	kworker/1:0H
19	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kdevtmpfs
20	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	netns
21	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_tasks_kthre
22	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kauditd
24	root	20	0	0	0	0	S	0.0	0.0	0:00.00	khungtaskd
25	root	20	0	0	0	0	S	0.0	0.0	0:00.00	oom_reaper
26	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	writeback
27	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kcompactd0
28	root	25	5	0	0	0	S	0.0	0.0	0:00.00	ksmd

Interactive Commands One of the powerful features of top is its interactive interface, which allows users to perform actions directly from the command. For example, users can kill a process by pressing **k** and entering the PID, or change the priority of a process using the **r** command. These interactive commands provide a convenient way to manage processes without needing to exit the top interface, making it a highly efficient tool for system administration.

Customizable Display The top command offers a customizable display, allowing users to tailor the information shown to their specific needs. Users can sort processes by different criteria, such as memory usage or CPU time, and choose which columns to display. Additionally, these preferences can be saved for future sessions, ensuring a consistent and personalized experience. This flexibility makes top a versatile tool that can be adapted to various monitoring and management tasks.

System Uptime and Load Average In addition to process information, top displays system uptime and load averages, which provide insights into the system's overall health

and stability. Uptime indicates how long the system has been running without a reboot, while load averages show the average system load over the past 1, 5, and 15 minutes. These metrics are useful for assessing the long-term performance and reliability of the system, helping users identify trends and potential issues.

All of these uses for the `top` command should sound similar to its Windows counterpart. Remembering those similarities is the best way to learn Linux.

Directory Navigation

In Linux, when you log in as root, your home directory is `/root`. For other users, the home directory is usually in the `/home` directory. For example, the home directory for a user logging in as *testuser* is `/home/testuser`. This information is stored in the `/etc/passwd` file. By default, only you have permission to save files in your home directory, and only you can create subdirectories in your home directory to organize your files further.

Linux supports the concept of a *current directory*, which is the directory on which all file and directory commands operate. After you log in, for example, your current directory is the home directory. To see the current directory, type the **pwd** command.

```
[user@localhost ~]$ pwd
/home/user
[user@localhost ~]$
```

To change the current directory, use the `cd` command. To change the current directory to `/usr/lib`, type the following:

```
[user@localhost ~]$ cd /usr/lib
[user@localhost lib]$
```

Then to change the directory to the `modules` subdirectory in `/usr/lib`, type the following command:

```
[user@localhost lib]$ cd modules
[user@localhost modules]$
```

You should notice that your command prompt reflects a portion of your current directory or focus directory. This functionality will differ depending on the shell. However, the Bash shell typically shows the focus directory. If you need the full path, we can use the `pwd` command, which will output `/usr/lib/modules` as the current full working directory path:

```
[user@localhost modules]$ pwd
/usr/lib/modules
[user@localhost modules]$
```

These two examples show that you can refer to a directory's name in two ways: with an absolute pathname or a relative pathname. An example of an absolute pathname is `/usr/lib`, which is an exact directory in the directory tree. (Think of the absolute pathname as the

complete mailing address for a package that the postal service will deliver to your next-door neighbor.) An example of a relative pathname is `modules`, which represents the `modules` subdirectory of the current directory, whatever that may be. (Think of the relative directory name as giving the postal carrier directions from your house to the one next door so that the carrier can deliver the package.)

If you type `cd modules` in `/usr/lib`, the current directory changes to `/usr/lib/modules`. However, if you type the same command in `/home/testuser`, the shell tries to change the current directory to `/home/testuser/modules`. This directory of `/home/testuser/modules` probably doesn't exist, but I think you get the idea.

Use the `cd` command without any arguments to change the current directory back to your *home directory*. No matter where you are, typing `cd` at the shell prompt brings you back home. The tilde character (`~`) is an alias that refers to your home directory. Thus, you can also change the current directory to your home directory by using the command `cd ~`. You can refer to another user's home directory by appending that user's name to the tilde. Thus, `cd ~comptia` changes the current directory to the home directory of the CompTIA user. Of course, we need to elevate our privileges with the `su -` command first, since the one user account will not have the privileges to another user's home directory. However, the root user does!

```
[user@localhost modules]$ cd
[user@localhost ~]$ pwd
/home/user
[user@localhost ~]$ su -
Password:
[root@localhost comptia]# whoami
root
[root@localhost ~]# cd ~comptia
[root@localhost comptia]# pwd
/home/comptia
[root@localhost comptia]#
```

A single dot (`.`) and two dots (`..`), often referred to as *dot-dot*, also have special meanings. A single dot (`.`) indicates the current directory, whereas two dots (`..`) indicate the parent directory. For example, if the current directory is `/usr/share`, you go one level up to `/usr` by typing the following:

```
[root@localhost comptia]# pwd
/home/comptia
[root@localhost comptia]# cd ..
[root@localhost home]# pwd
/home
[root@localhost home]# ls -la
total 0
drwxr-xr-x.  4 root   root   33 Jan  1 13:44 .
```

```
dr-xr-xr-x. 18 root    root    235 Jan  1 12:11 ..
drwx-----.  2 comptia comptia  62 Jan  1 13:44 comptia
drwx-----.  3 user     user    114 Jan  1 13:22 user
[root@localhost home]#
```

After executing the `ls -la` command, you can see the directory structure you just navigated. You can also see the single dot and double dot referenced in the directory listing. In the next section, we'll examine directory listing more closely.

Directory Listings

You can get a directory listing by using the `ls` command. By default, the `ls` command, without any options, displays the contents of the current directory in a compact, multi-column format. To tell the directories and files apart, use the `-F` option (`ls -F`). The output will show the directory names with a slash (/) appended to them. Plain filenames appear as is. The at sign (@) appended to a listing indicates that this file is a link to another file. (In other words, this filename simply refers to another file; it's a shortcut.) An asterisk (*) is appended to executable files. (The shell can run any executable file.)

```
[user@localhost /]$ ls -F
afs/  boot/  etc/  lib@   media/  opt/  root/  sbin@  sys/  usr/
bin@  dev/  home/  lib64@ mnt/    proc/  run/  srv/  tmp/  var/
[user@localhost /]$
```

You can see even more detailed information about the files and directories with the `-l` (long format) option. The rightmost column shows the name of the directory entry. The date and time before the name show when the last modifications to that file were made. To the left of the date and time is the size of the file in bytes. The file's group and owner appear to the left of the column that shows the file size. The next number to the left indicates the number of links to the file. (A *link* is like a shortcut in Windows.)

```
[user@localhost /]$ ls -l
total 20
dr-xr-xr-x.  2 root root    6 Jul 16 20:00 afs
lrwxrwxrwx.  1 root root    7 Jul 16 20:00 bin -> usr/bin
dr-xr-xr-x.  5 root root 4096 Jan  1 12:31 boot
drwxr-xr-x. 20 root root 3920 Jan  1 12:28 dev
drwxr-xr-x. 108 root root 8192 Jan  1 13:44 etc
drwxr-xr-x.  4 root root   33 Jan  1 13:44 home
lrwxrwxrwx.  1 root root    7 Jul 16 20:00 lib -> usr/lib
lrwxrwxrwx.  1 root root    9 Jul 16 20:00 lib64 -> usr/lib64
drwxr-xr-x.  2 root root    6 Jul 16 20:00 media
drwxr-xr-x.  3 root root   18 Jan  1 12:12 mnt
```


[Output Cut]

```
[user@localhost ~]$
```

Linux Permissions

The leftmost column of a long style directory listing shows the file's permission settings, which determine who can read, write, or execute the file. This column shows a sequence of nine characters, which appear as `rw-rw-rw-` when each letter is present. Each letter indicates a specific permission. A hyphen (-) in place of a letter indicates no permission for a specific operation on the file. Think of these nine letters as three groups of three letters (`rw-`), interpreted as follows:

Leftmost Group Controls the read, write, and execute permissions of the file's owner. In other words, if you see `rw-` in this position, the file's owner can read (r), write (w), and execute (x) the file. A hyphen in the place of a letter indicates no permission. Thus, the string `rw-` means that the owner has read and write permissions but not execute permission. Although executable programs (including shell programs) typically have execute permission, directories treat execute permission as equivalent to *use* permission: a user must have execute permission on a directory before they can open and read the contents of the directory.

Middle Group Controls the read, write, and execute permissions of any user belonging to that file's group.

Rightmost Group Controls the read, write, and execute permissions of all other users (collectively thought of as *the world*).

Thus, a file with the permission setting `rw-----` is accessible only to the file's owner, whereas the permission setting `rw-r--r--` makes the file readable by the world.



An interesting feature of the `ls` command is that it doesn't list any file the name of which begins with a period. To see these files, you must use the `ls` command with the `-a` option, as follows:

```
[user@localhost ~]$ ls -a
.  .bash_history  .bash_profile  .config
.. .bash_logout  .bashrc        .lessht
[user@localhost ~]$
```

Most Linux commands take single-character options, each with a hyphen as a prefix. When you want to use several options, type a hyphen and concatenate (string together) the option letters, one after another. Thus, `ls -al` is equivalent to `ls -a -l` as well as to `ls -l -a`.

```
[user@localhost ~]$ ls -al
```

```
total 20
drwx-----. 3 user user 114 Jan  1 13:22 .
drwxr-xr-x. 4 root root  33 Jan  1 13:44 ..
-rw-----. 1 user user   7 Jan  1 12:16 .bash_history
-rw-r--r--. 1 user user  18 Aug 11 20:00 .bash_logout
-rw-r--r--. 1 user user 144 Aug 11 20:00 .bash_profile
-rw-r--r--. 1 user user 522 Aug 11 20:00 .bashrc
drwx-----. 3 user user  20 Jan  1 12:47 .config
-rw-----. 1 user user  20 Jan  1 13:22 .lessht
[user@localhost ~]$
```

Changing Permissions and Ownership

You may need to change a file's permission settings to protect it from others. Use the `chmod` command to change the permission settings of a file or a directory. To use `chmod` effectively, you have to specify the permission settings. A good way is to concatenate letters from the columns of Table 4.2 in the order shown (who/action/permission). You use only the single character from each column—the text in parentheses is for explanation only.

TABLE 4.2 Letter codes for file permissions

Who	Action	Permission
u (user)	+ (add)	r (read)
g (group)	– (remove)	w (write)
o (others)	= (assign)	x (execute)
a (all)	s (set user ID)	

For example, to give everyone read access to all the files in a directory, pick a (for *all*) from the first column, + (for *add*) from the second column, and r (for *read*) from the third column, to come up with the permission setting `a+r`. Then use the set of options with `chmod`, as follows:

```
[user@localhost folder]$ ls -l
total 0
-rw-r--r--. 1 user user   0 Jan  1 15:07 bin.exe
-rw-r--r--. 1 user user  34 Jan  1 15:05 file1.txt
-rw-r--r--. 1 user user  34 Jan  1 15:05 file2.txt
-rw-r--r--. 1 user user  36 Jan  1 15:05 file3.txt
```

```
[user@localhost folder]$ chmod a+rw *
[user@localhost folder]$ ls -l
total 0
-rw-rw-rw-. 1 user user  0 Jan  1 15:07 bin.exe
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan  1 15:05 file3.txt
[user@localhost folder]$
```

To permit everyone to execute one specific file, type the following:

```
[user@localhost folder]$ chmod a+x bin.exe
[user@localhost folder]$ ls -l
total 0
-rwxrwxrwx. 1 user user  0 Jan  1 15:07 bin.exe
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan  1 15:05 file3.txt
[user@localhost folder]$
```

Sometimes you have to change a file's user or group ownership in order for everything to work correctly. For example, suppose you're instructed to create a directory named *cups* and give it the ownership of user ID *lp* and group ID *sys*. You create the cups directory with the command `mkdir` as follows:

```
[user@localhost folder]$ mkdir cups
[user@localhost folder]$ ls -l
total 0
-rwxrwxrwx. 1 user user  0 Jan  1 15:07 bin.exe
drwxr-xr-x. 2 user user  6 Jan  1 15:13 cups
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan  1 15:05 file3.txt
[user@localhost folder]$
```

If you check the file's details with the `ls -l` command, you see that the user and group ownership are both assigned to user. To change the owner, use the `chown` command. For example, to change the ownership of the cups directory to user ID *lp* and group ID *sys*, type the following:

```
[user@localhost folder]$ sudo chown lp:sys cups
[sudo] password for user:
[user@localhost folder]$ ls -l
total 0
-rwxrwxrwx. 1 user user  0 Jan  1 15:07 bin.exe
```

```
drwxr-xr-x. 2 lp sys 6 Jan 1 15:13 cups
-rw-rw-rw-. 1 user user 34 Jan 1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan 1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan 1 15:05 file3.txt
[user@localhost folder]$
```

Working with Files

Working with files in Linux is a fundamental skill for anyone using the Linux operating system. Whether you're a beginner or an experienced user, mastering file operations is essential for efficient system management and productivity. In the following, we'll explore key file management tasks such as copying files, moving files, and viewing file contents. Each of these operations plays a crucial role in organizing and manipulating data, ensuring that you can handle your files with ease and precision. Let's dive into the world of Linux file management and discover how to perform these tasks effectively.

Copying Files

To copy files from one directory to another, use the `cp` command. If you want to copy a file to the current directory but retain the original name, use a period (.) as the second argument of the `cp` command. Thus, the following command copies the `host.conf` file from the `/etc` directory to the current directory (denoted by a single period):

```
[user@localhost ~]$ cp /etc/host.conf .
[user@localhost ~]$ ls -l
total 4
drwxr-xr-x. 3 user user 84 Jan 1 15:13 folder
-rw-r--r--. 1 user user 9 Jan 1 15:22 host.conf
[user@localhost ~]$
```

The `cp` command makes a new copy of a file and leaves the original intact.

If you want to copy the entire contents of a directory—including all subdirectories and their contents—to another directory, use the command `cp -ar sourcedir destdir`. (This command copies everything in the *sourcedir* directory to the *destdir* directory.) For example, to copy all the files from the `/etc/ssl` directory to the current directory, type the following command:

```
[user@localhost ~]$ cp -ar /etc/ssl .
[user@localhost ~]$ ls -l
total 4
drwxr-xr-x. 3 user user 84 Jan 1 15:13 folder
-rw-r--r--. 1 user user 9 Jan 1 15:22 host.conf
drwxr-xr-x. 2 user user 91 Jan 1 12:11 ssl
[user@localhost ~]$
```

Moving Files

To move a file to a new location, use the `mv` command; the syntax is similar to the `cp` command, such as `mv source destination`. The original file or folder will be gone, and a new copy appears at the destination. You can use `mv` to rename a file. If you want to change the name of `host.conf` to `host.conf.bak`, use the `mv` command as follows:

```
[user@localhost ~]$ ls -l
total 4
drwxr-xr-x. 3 user user 84 Jan  1 15:29 folder
-rw-r--r--. 1 user user  9 Jan  1 15:22 host.conf
drwxr-xr-x. 2 user user 91 Jan  1 12:11 ssl
[user@localhost ~]$ mv host.conf host.conf.bak
[user@localhost ~]$ ls -l
total 4
drwxr-xr-x. 3 user user 84 Jan  1 15:29 folder
-rw-r--r--. 1 user user  9 Jan  1 15:22 host.conf.bak
drwxr-xr-x. 2 user user 91 Jan  1 12:11 ssl
[user@localhost ~]$
```

If you want to move the `host.conf.bak` file to a subdirectory named `folder`, use the following command:

```
[user@localhost ~]$ mv host.conf.bak folder/
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 3 user user 105 Jan  1 15:31 folder
drwxr-xr-x. 2 user user  91 Jan  1 12:11 ssl
[user@localhost ~]$ ls -l folder/
total 4
-rwxrwxrwx. 1 user user  0 Jan  1 15:07 bin.exe
drwxr-xr-x. 2 lp  sys  6 Jan  1 15:13 cups
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan  1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan  1 15:05 file3.txt
-rw-r--r--. 1 user user  9 Jan  1 15:22 host.conf.bak
[user@localhost ~]$
```

An interesting feature of `mv` is that you can use it to move entire directories (with all their subdirectories and files) to a new location. If you have a directory named `data` that contains many files and subdirectories, you can move that entire directory structure to `old_data` by using the following command:

```
[user@localhost folder]$ ls -l
total 4
```

```

-rwxrwxrwx. 1 user user 0 Jan 1 15:07 bin.exe
drwxr-xr-x. 2 lp sys 6 Jan 1 15:13 cups
drwxr-xr-x. 2 user user 57 Jan 1 15:34 files
-rw-r--r--. 1 user user 9 Jan 1 15:22 host.conf.bak
[user@localhost folder]$ ls -l files
total 0
-rw-rw-rw-. 1 user user 34 Jan 1 15:05 file1.txt
-rw-rw-rw-. 1 user user 34 Jan 1 15:05 file2.txt
-rw-rw-rw-. 1 user user 36 Jan 1 15:05 file3.txt
[user@localhost folder]$ mv files ~/
[user@localhost folder]$ cd ~
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 57 Jan 1 15:34 files
drwxr-xr-x. 3 user user 54 Jan 1 15:35 folder
drwxr-xr-x. 2 user user 91 Jan 1 12:11 ssl
[user@localhost ~]$

```

To delete files, use the `rm` command. For example, to delete a file named `file4.txt`, type the following command:

```

[user@localhost ~]$ ls -l
total 0
-rw-r--r--. 1 user user 34 Jan 1 15:36 file4.txt
drwxr-xr-x. 2 user user 57 Jan 1 15:34 files
drwxr-xr-x. 3 user user 54 Jan 1 15:35 folder
drwxr-xr-x. 2 user user 91 Jan 1 12:11 ssl
[user@localhost ~]$ rm file4.txt
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 57 Jan 1 15:34 files
drwxr-xr-x. 3 user user 54 Jan 1 15:35 folder
drwxr-xr-x. 2 user user 91 Jan 1 12:11 ssl
[user@localhost ~]$

```

To delete a folder, you need to add the `-rf` switch to the command line. The `-r` switch is for recursive operations and the `-f` is for force, so you don't get a bunch of message prompts asking if it is okay to delete a subfolder. The following displays what happens if you forget the `-rf` on the command line:

```

[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 57 Jan 1 15:34 files

```

```
drwxr-xr-x. 3 user user 54 Jan  1 15:35 folder
drwxr-xr-x. 2 user user 91 Jan  1 12:11 ssl
[user@localhost ~]$ rm folder
rm: cannot remove 'folder': Is a directory
[user@localhost ~]$ rm -rf folder
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 57 Jan  1 15:34 files
drwxr-xr-x. 2 user user 91 Jan  1 12:11 ssl
[user@localhost ~]$
```



Be careful with the `rm` command, especially when you log in as root. You can inadvertently delete important files with `rm`.

Viewing Files

Sometime you just need to view the contents of a file. To view the contents of a file, use the `cat` command. The `cat` command will display the contents of a file, as shown in the following example:

```
[user@localhost ~]$ cat /etc/hosts
# Loopback entries; do not change.
# For historical reasons, localhost precedes localhost.localdomain:
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
# See hosts(5) for proper format and other examples:
# 192.168.1.10 foo.example.org foo
# 192.168.1.13 bar.example.org bar
[user@localhost ~]$
```


The `cat` file can also be used to concatenate multiple files. For example, you may have two files that you want to join together. By using the `>` redirector, you can redirect the output of two or more files into one file, as follows:

```
[user@localhost files]$ cat file1.txt
This is the contents of file one!
[user@localhost files]$ cat file2.txt
This is the contents of file two!
[user@localhost files]$ cat file2.txt file1.txt > file3.txt
[user@localhost files]$ cat file3.txt
This is the contents of file two!
This is the contents of file one!
[user@localhost files]$
```

Editing Files

The original text-based editor that came with Linux/UNIX was the `vi` editor, which was short for *visual editor*. The original `vi` editor was not end user-friendly, and it often required a cheat sheet to get anything done. Thankfully, the `nano` and `pico` editors were adopted by Linux, which made editing text-based config files much easier. You can launch the `nano` editor by typing the **nano** command followed by the file you want to edit, such as `nano file3`. This will launch the editor, as shown in Figure 4.24.

FIGURE 4.24 The nano editor



Once the editor is launched, it acts similar to the Windows Notepad utility. You can use the arrow keys to navigate text; find text; replace text; and cut, copy, and paste text. The editor is also very intuitive; by using the Control key sequences, you can perform all the functions and even get help.

Locating Files and Content

One of the most important functions you can perform on files is finding them. In Windows we can use the familiar `Ctrl+F` key sequence to search for files in a folder structure. In Linux the `find` command allows you to do the same. You can find a file in a folder structure, as shown here:

```
[user@localhost dev]$ find ttyS12
ttyS12
[user@localhost dev]$
```


The `grep` command in Linux is a powerful utility used for searching and filtering text within files. Whether you're a system administrator, developer, or just a casual user, mastering `grep` can significantly enhance your ability to find and manipulate data efficiently. The `grep` command includes a method of using regular expressions, and by combining `grep` with other commands, you'll get more advanced text processing.

The primary use of `grep` is to search for specific patterns within files. For example, to find all occurrences of the word `Error` in a log file, you can use the following command:

```
[root@localhost ~]# grep "Error" /var/log/messages
Jan  1 12:16:26 localhost kernel: RAS: Correctable Errors collector
initialized.
Jan  1 12:28:33 localhost kernel: RAS: Correctable Errors collector
initialized.
[root@localhost ~]#
```

The `grep` command also supports regular expressions, allowing for more complex and flexible search patterns. For instance, to find lines that contain either `error` or `info` you can use the command `grep -E "Error|Info" /var/log/messages`. The `-E` option enables extended regular expressions, providing powerful pattern-matching capabilities.

```
[root@localhost ~]# grep -E "Error|Info" /var/log/messages
Jan  1 12:16:26 localhost kernel: RAS: Correctable Errors collector
initialized.
Jan  1 12:16:29 localhost systemd[1]: systemd-hibernate-clear.service - Clear
Stale Hibernate Storage Info was skipped because of an unmet condition check
(ConditionPathExists=/sys/firmware/efi/efivars/HibernateLocation-8cf2644b-4b0b
-428f-9387-6d876050dc67).
Jan  1 12:28:33 localhost kernel: RAS: Correctable Errors collector
initialized.
Jan  1 12:28:36 localhost systemd[1]: systemd-hibernate-clear.service - Clear
Stale Hibernate Storage Info was skipped because of an unmet condition check
(ConditionPathExists=/sys/firmware/efi/efivars/HibernateLocation-8cf2644b-4b0b
-428f-9387-6d876050dc67).
[root@localhost ~]#
```

Sometimes, you may need to perform case-insensitive searches. The `-i` option allows you to ignore case differences. The command line `grep -i -E "Error|Info" /var/log/messages` will search for `error` or `info`, as shown here:

```
[root@localhost ~]# grep -i -E "Error|Info" /var/log/messages
Jan  1 12:16:26 localhost kernel: ACPI: Using ACPI (MADT) for SMP
configuration information
Jan  1 12:16:26 localhost kernel: RAS: Correctable Errors collector
initialized.
Jan  1 12:16:29 localhost systemd[1]: systemd-hibernate-clear.service - Clear
Stale Hibernate Storage Info was skipped because of an unmet condition check
(ConditionPathExists=/sys/firmware/efi/efivars/HibernateLocation-8cf2644b-4b0b
-428f-9387-6d876050dc67).
```

```
Jan 1 12:16:30 localhost mcelog[934]: mcelog: ERROR: AMD Processor family 23:
mcelog does not support this processor. Please use the edac_mce_amd module
instead.
```

```
Jan 1 12:16:30 localhost rsyslogd[936]: [origin software="rsyslogd"
swVersion="8.2312.0-5.fc41" x-pid="936" x-info="https://www.rsyslog
.com"] start
```

```
Jan 1 12:16:31 localhost chronyd[970]: Using leap second list /usr/share/
zoneinfo/leap-seconds.list
```

```
Jan 1 12:16:31 localhost NetworkManager[1057]: <info> [1735751791.6809]
NetworkManager (version 1.50.0-1.fc41) is starting... (boot:c0e22ce0-c23f-45ee
-9801-0953d63c8548)
```

```
Jan 1 12:16:31 localhost NetworkManager[1057]: <info> [1735751791.6820] Read
config: /etc/NetworkManager/NetworkManager.conf (lib: 22-wifi-mac-addr.conf,
99-nvme-nbft-no-ignore-carrier.conf)
```

```
[Output Cut]
```

```
[root@localhost ~]#
```

You can even search files recursively for a pattern within a directory and its subdirectories by using the `-r` option. The command `grep -r "sshd" /var/log/` will search for `sshd` in all files under the `/var/log/` directory, providing a search across multiple files, as shown here:

```
[root@localhost ~]# grep -r "sshd" /var/log/
/var/log/anaconda/journal.log:Jan 01 17:12:25 localhost.localdomain
groupadd[3823]: group added to /etc/group: name=sshd, GID=74
/var/log/anaconda/journal.log:Jan 01 17:12:25 localhost.localdomain
groupadd[3823]: group added to /etc/gshadow: name=sshd
/var/log/anaconda/journal.log:Jan 01 17:12:25 localhost.localdomain
groupadd[3823]: new group: name=sshd, GID=74
/var/log/boot.log:[ OK ] Created slice system-sshd\x2dkeygen.slice - Slice /
system/ssh-keygen.
/var/log/boot.log:[ OK ] Listening on sshd-UNIX-local.socket - OpenSSH
Server Socket (systemd-ssh-generator, AF_UNIX Local).
/var/log/boot.log:[ OK ] Listening on sshd-vsock.socket - OpenSSH Server
Socket (systemd-ssh-generator, AF_VSOCK).
/var/log/boot.log: Starting sshd-keygen@ecdsa.service - OpenSSH ecdsa
Server Key Generation...
/var/log/boot.log: Starting sshd-keygen@ed25519.service - OpenSSH
ed25519 Server Key Generation...
/var/log/boot.log: Starting sshd-keygen@rsa.service - OpenSSH rsa
Server Key Generation...
```

```
[Output Cut]
```

```
[root@localhost ~]#
```

The `grep` command can be combined with other commands using pipes to create powerful command-line workflows. For example, to list all running processes and filter for those containing `ssh` you can use the following:

```
[root@localhost ~]# ps aux | grep "ssh"
root      1215  0.0  0.2 15416 11076 ?        Ss   12:28   0:00 sshd-
session: user [priv]
user      1232  0.0  0.1 15680  7080 ?        S    12:28   0:01 sshd-
session: user@pts/0
root      11744 0.0  0.1  8160  6024 ?        Ss   12:31   0:00 sshd: /usr/
sbin/sshd -D [listener] 0 of 10-100 startups
root      14035 0.0  0.0  6500  2136 pts/0    S+   20:21   0:00 grep
--color=auto ssh
[root@localhost ~]#
```

By mastering these `grep` techniques, you can efficiently search and filter text in Linux, making `grep` an indispensable tool for managing and analyzing data.

In Exercise 4.3 you will work with some basic files on the Linux operating system to get some exposure to working with files.

EXERCISE 4.3

Working with Files

1. Log into a Linux or macOS operating system and open a terminal.
2. Type `cd ~` and press Enter. This will change the directory to your home folder.
3. Type `cat > file1` and press Enter. Begin typing **this is file1 contents**, then press Control+D.
4. Type `cat > file2` and press Enter. Begin typing **this is file2 contents**, then press Control+D.
5. Type `ls` and press Enter to examine the two files you've created.
6. Type `cat file2 file1 > file3` and press Enter. This will concatenate the two files.
7. Type `cat file3` and press Enter to examine the contents of `file3`.

In this example, we used the `cat` command to redirect console input to a file, such as `file1` and `file2`. (We could have used the `nano` command if it was available to achieve the same outcome.) We then used the `cat` command to concatenate the two files.

Working with Directories

To organize files in your home directory, you have to create new directories. Use the `mkdir` command to create a directory. For example, to create a directory named `images` in the current directory, type the following:

```
[user@localhost ~]$ mkdir images
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 6 Jan  1 20:24 images
[user@localhost ~]$
```

After you create the directory, you can use the `cd images` command to change to that directory.

```
[user@localhost ~]$ cd images
[user@localhost images]$
```

You can create an entire directory tree by using the `-p` option with the `mkdir` command. For example, suppose you need to create a directory path of `files/src/jscript/examples`. To create this directory hierarchy, type the following command:

```
[user@localhost ~]$ mkdir -p files/src/jscript/examples
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 3 user user 17 Jan  1 20:27 files
drwxr-xr-x. 2 user user  6 Jan  1 20:24 images
[user@localhost ~]$ cd files
[user@localhost files]$ ls -l
total 0
drwxr-xr-x. 3 user user 21 Jan  1 20:27 src
[user@localhost files]$ cd src/jscript/examples/
[user@localhost examples]$
```

When you no longer need a directory, use the `rmdir` command to delete it. You can delete a directory only when the directory is empty. To remove an empty directory tree, you can use the `-p` option, as follows:

```
[user@localhost examples]$ cd ~
[user@localhost ~]$ rmdir -p files/src/jscript/examples/
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 6 Jan  1 20:24 images
[user@localhost ~]$
```

This command removes the empty parent directories of applets. The command stops when it encounters a directory that's not empty.

Directory Usage

The `du` (disk usage) command in Linux is an invaluable tool for assessing and managing disk space usage. Understanding how to use `du` command can help you find out-of-control disk usage. In the following, we examine various applications of the `du` command, including checking directory sizes and summarizing disk usage.

The most common use of the `du` command is to check the size of directories. For example, to display the size of the `/home` directory, you can use the command `sudo du -sh /home`. The `-s` option provides a summary of the directory size, and the `-h` option makes the output human-readable, showing sizes in KB, MB, or GB. Of course, we need to use the `sudo` command, since our regular user account will not have permissions to everyone's home folder:

```
[user@localhost ~]$ sudo du -sh /home
[sudo] password for user:
32K      /home
[user@localhost ~]$
```

If we want to get a summary of disk usage for all directories within a specific path, you can add the asterisk to the end of the path. This command will display the size of each directory and file in the current directory, providing a quick overview of disk usage:

```
[user@localhost ~]$ sudo du -sh /home/*
12K      /home/comptia
20K      /home/user
[user@localhost ~]$
### Displaying Disk Usage for All Files
```

If you need to see the disk usage for all files within a directory, including subdirectories, use the command `sudo du -ah /home`. The `-a` option includes all files, and the `-h` option makes the output human-readable.

```
[user@localhost ~]$ sudo du -ah /home
4.0K     /home/user/.bash_logout
4.0K     /home/user/.bash_profile
4.0K     /home/user/.bashrc
4.0K     /home/user/.bash_history
0        /home/user/.config/procps
0        /home/user/.config
4.0K     /home/user/.lessht
0        /home/user/images
20K      /home/user
4.0K     /home/comptia/.bash_logout
4.0K     /home/comptia/.bash_profile
4.0K     /home/comptia/.bashrc
```

```
12K      /home/comptia
32K      /home/
[user@localhost ~]$
```

By mastering these `du` techniques, you can efficiently manage disk space in Linux, ensuring that your filesystem remains organized and optimized.

Networking Utilities

Networking utilities in Linux are essential tools for managing and troubleshooting network connections. Understanding these utilities can greatly enhance your ability to diagnose and resolve network issues.

Managing Network Interfaces

Just as you can use the `ipconfig` command to see the status of IP configuration with Windows, the `ifconfig` command can be used in Linux. You can get information about the usage of the `ifconfig` command by using `ifconfig -help`. The following output provides an example of the basic `ifconfig` command run on a Linux system:

```
[user@localhost ~]$ ifconfig
ens160: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
    inet 172.16.1.13  netmask 255.240.0.0  broadcast 172.31.255.255
    inet6 fe80::20c:29ff:fe01:5345  prefixlen 64  scopeid 0x20<link>
    ether 00:0c:29:01:53:45  txqueuelen 1000  (Ethernet)
    RX packets 464139  bytes 624240392 (595.3 MiB)
    RX errors 0  dropped 32191  overruns 0  frame 0
    TX packets 79223  bytes 7326946 (6.9 MiB)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
    inet 127.0.0.1  netmask 255.0.0.0
    inet6 ::1  prefixlen 128  scopeid 0x10<host>
    loop txqueuelen 1000  (Local Loopback)
    RX packets 4  bytes 240 (240.0 B)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 4  bytes 240 (240.0 B)
    TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

```
[user@localhost ~]$
```

In addition to using `ifconfig`, Linux users can use the `iwconfig` command to view the state of their wireless network. By using `iwconfig`, you can view such important information as the link quality, access point (AP) MAC address, data rate, and encryption keys, which can be helpful in ensuring that the parameters in the network are consistent.

The `ip` utility has replaced the `ifconfig` utility on most modern Linux distributions. Red Hat Enterprise Linux has adopted the `ip` utility, and it's similar to the `ifconfig` utility. You can configure an address or show the configured IP addresses similar to the `ifconfig` utility. The output even looks very familiar to the `ifconfig` utility:

```
[root@localhost ~]# ip addr add 172.16.1.200/12 eth0
[root@localhost ~]# ip addr
eth0: <BROADCAST, MULTICAST, UP> mtu 1500 qlen 1000
    link/ether 00:0c:29:e9:08:92 brd ff:ff:ff:ff:ff:ff:ff:ff
    inet 172.16.1.200/12 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::20c:29ff:fee9:892/64 scope link
        valid_lft forever preferred_lft forever
lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 noqueue
    link/loopback 00:00:00:00:00:00:00:00 brd 00:00:00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
[root@localhost ~]#
```

Resolvers

Resolvers in Linux are essential components that translate human-readable domain names into IP addresses, enabling communication over the Internet. This process, known as domain name resolution, is handled by the resolver library, which consists of functions that query DNS servers to obtain the corresponding IP addresses. When a user types a URL into their browser, the resolver takes over, converting the domain name into an IP address that the system can use to establish a connection. This seamless translation is crucial for accessing websites, sending emails, and performing other network-related tasks.

The `/etc/resolv.conf` file is how you configure the resolver for the operating system in Linux. This file contains information about DNS servers that the system should use for resolving domain names. Typically, it includes one or more `nameserver` entries, each specifying the IP address of a DNS server. For example, a basic `/etc/resolv.conf` file might look like this:

```
[user@localhost ~]$ cat /etc/resolv.conf

nameserver 8.8.8.8
nameserver 8.8.4.4
search sybex.com

[user@localhost ~]$
```

In this example, the resolver will first attempt to query the DNS server at 8.8.8.8. If that server is unavailable, it will try the next one at 8.8.4.4. This configuration ensures that the system can resolve domain names even if one of the DNS servers fails. Additionally, the `/etc/resolv.conf` file can include other options, such as search directives, which specify default search domains to append to queries. In the previous example, `sybex.com` will be appended to queries.

Modern Linux distributions often use tools like `systemd-resolved` to dynamically manage the contents of `/etc/resolv.conf`, ensuring that the resolver configuration adapts to changing network conditions. For instance, when the device connects to a new network or a VPN, `systemd-resolved` can dynamically adjust the DNS server entries to reflect the new network's DNS settings. This ensures that the system always uses the most appropriate DNS servers without requiring manual intervention. When `system-resolved` is employed, the configuration file will typically be `/etc/systemd/resolved.conf` and it has several different options, beyond the scope of this exam.

Resolving Domain Names

The Domain Information Groper (`dig`) tool is almost identical to the Windows `nslookup` tool and has become an adopted standard for name resolution testing on Linux/UNIX operating systems. The tool allows you to resolve any resource record for a given host and direct the query to a specific server.

The command does not offer an interactive mode like the `nslookup` command. The command by default queries A records for a given host, and the output has debugging turned on by default.

In the following example, you see a query being performed on the DNS server of 8.8.8.8 for an MX record of `sybex.com`. The debugging output shows that one query was given, two answers were retrieved, and nothing was authoritative (not the primary servers). The output also details the query made and the answers returned:

```
[root@localhost ~]# dig @8.8.8.8 mx sybex.com

; <<>> DiG 9.9.5-3ubuntu0.13-Ubuntu <<>> @8.8.8.8 mx sybex.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 49694
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 512
;; QUESTION SECTION:
;sybex.com.                IN      MX

;; ANSWER SECTION:
```



```
sybex.com.      899  IN   MX   10 cluster1.us.messagelabs.com.  
sybex.com.      899  IN   MX   20 cluster1a.us.messagelabs.com.
```

```
;; Query time: 76 msec  
;; SERVER: 8.8.8.8#53(8.8.8.8)  
;; WHEN: Wed Nov 01 21:43:32 EDT 2017  
;; MSG SIZE rcvd: 104
```

```
[root@localhost ~]#
```

The HOSTS File

The `/etc/hosts` file has a history that dates back to the early days of the Internet. Before the Domain Name System (DNS) was developed in the 1980s, the `HOSTS.TXT` file was used to map hostnames to IP addresses. This file was manually maintained and distributed among ARPANET members by the Stanford Research Institute. As the Internet grew, the manual maintenance of this file became impractical, leading to the creation of DNS. Despite the advent of DNS, the `/etc/hosts` file remains an integral part of modern operating systems, providing a simple and effective way to manage hostname resolution locally for the operating system.

Today the `/etc/hosts` file is primarily used for local hostname resolution. It allows systems to map IP addresses to hostnames without the need to query a DNS server. This can be particularly useful in development and testing environments, where developers need to simulate how websites will behave on a live server. By mapping domain names to local IP addresses, developers can test their applications in a controlled environment. Additionally, the `/etc/hosts` file can be used to override DNS settings for specific hostnames, providing a quick and easy way to redirect traffic or block access to certain websites by mapping them to non-routable IP addresses.

Administrators often edit the `/etc/hosts` file for various reasons. One common use is to troubleshoot network issues by temporarily overriding DNS settings to ensure that a hostname resolves to the correct IP address. This can help identify and resolve DNS-related problems without waiting for DNS propagation. Another use is during server migrations, where administrators can map old domain names to new IP addresses to ensure a smooth transition with minimal downtime. By updating the `/etc/hosts` file, administrators can test new server configurations and verify that everything is working correctly before making DNS changes.

Network Connectivity

Network connectivity is a fundamental aspect of Linux systems, enabling communication between devices and networks. This section will cover the tools that CompTIA has identified as exam objectives. However, there are many other tools depending on what you are trying to diagnose.

The *ping* Command

The *ping* command is a fundamental network utility used to test the reachability of a host on an IP network and measure the round-trip time for messages sent from the originating host to a destination computer. It operates by sending Internet Control Message Protocol (ICMP) Echo Request packets to the target host and waiting for an Echo Reply. This process helps in diagnosing network connectivity issues, determining network latency, and verifying that a specific IP address is active. For example, to check if a server with the IP address 8.8.8.8 is reachable, you can use the following command:

```
[user@localhost ~]$ ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=114 time=17.5 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=114 time=30.0 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=114 time=23.9 ms
^C
--- 8.8.8.8 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 17.505/23.794/29.973/5.090 ms
[user@localhost ~]$
```

This command will send a series of packets to the specified IP address and display the response time for each packet. The ping of packet will continue indefinitely. To stop the constant ping, you need to press Control+C to stop the command.

The *traceroute* Command

The *traceroute* command, similar to the *tracert* Windows command, is another valuable network diagnostic tool used to trace the path that packets take from your computer to a specified destination. By sending a series of ICMP or UDP packets with incrementally increasing Time-To-Live (TTL) values, *traceroute* reveals each hop along the route to the target host. This helps in identifying where delays or failures occur within the network. It also helps with troubleshooting connectivity issues and understanding network topology. For instance, to trace the route to 8.8.8.8, you can use the following command:

```
[user@localhost ~]$ traceroute 8.8.8.8
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 60 byte packets
 1 _gateway (172.16.1.1) 0.977 ms 0.867 ms 0.811 ms
 2 100.93.196.195 (100.93.196.195) 15.302 ms 100.93.196.194 (100.93.196.194)
15.351 ms 15.421 ms
 3 po-318-346-rur402.pittsburgh.pa.pitt.comcast.net (68.86.100.233) 15.303
ms 15.417 ms 15.349 ms
 4 po-400-xar02.pittsburgh.pa.pitt.comcast.net (68.85.75.65) 24.003 ms po-2-
rur402.pittsburgh.pa.pitt.comcast.net (96.108.91.26) 23.918 ms po-400-
xar02.pittsburgh.pa.pitt.comcast.net (68.85.75.65) 23.993 ms
```

```

5  be-22-ar01.mckeesport.pa.pitt.comcast.net (162.151.65.5) 25.465
ms 25.520
   ms po-400-xar02.pittsburgh.pa.pitt.comcast.net (68.85.75.65) 23.944 ms
6  be-31611-cs01.pittsburgh.pa.ibone.comcast.net (96.110.42.161)
25.676 ms be-
31621-cs02.pittsburgh.pa.ibone.comcast.net (96.110.42.165) 24.214 ms
be-31611-
cs01.pittsburgh.pa.ibone.comcast.net (96.110.42.161) 24.258 ms
7  be-31621-cs02.pittsburgh.pa.ibone.comcast.net (96.110.42.165)
24.134 ms be-
1311-cr11.pittsburgh.pa.ibone.comcast.net (96.110.38.138) 17.068 ms
16.090 ms
8  be-302-cr12.beaumeade.va.ibone.comcast.net (96.110.32.101) 21.190
ms be-301-cr12.beaumeade.va.ibone.comcast.net (96.110.39.165) 21.161 ms
21.115 ms
9  be-1212-cs02.beaumeade.va.ibone.comcast.net (68.86.84.149) 33.702 ms
24.445 ms be-3212-pe12.ashburn.va.ibone.comcast.net (96.110.34.118)
32.619 ms
10 be-3112-pe12.ashburn.va.ibone.comcast.net (96.110.34.114) 24.341 ms be-
3312-pe12.ashburn.va.ibone.comcast.net (96.110.34.122) 33.054 ms be-3213-
pe13.ashburn.va.ibone.comcast.net (68.86.166.214) 32.922 ms
11 * * *
12 * * *
13 dns.google (8.8.8.8) 21.258 ms 21.758 ms 142.251.252.71 (142.251.252.71)
22.597 ms
[user@localhost ~]$

```

This command will display each hop along the path to 8.8.8.8, including the IP address and response time for each intermediate router, providing a detailed map of the network route and highlighting any potential points of failure or latency.

The *curl* Command

The *curl* command is a versatile and powerful tool used for transferring data to and from a server using various protocols, including HTTP, HTTPS, FTP, and more. It is commonly used for testing and interacting with web services, downloading files, and performing network requests. With *curl*, you can send GET, POST, and other types of HTTP requests, making it an essential utility for developers and system administrators. For example, to fetch the contents of a web page, you can use the following command:

```

[user@localhost ~]$ curl https://8.8.8.8
<HTML><HEAD><meta http-equiv="content-type" content="text/html; charset=utf-8">
<TITLE>302 Moved</TITLE></HEAD><BODY>
<H1>302 Moved</H1>
The document has moved

```

```
<A HREF="https://dns.google/">here</A>.  
</BODY></HTML>  
[user@localhost ~]$
```

This command will retrieve the HTML content of the specified URL and display it in the terminal, allowing you to inspect the response from the server. If you wanted to capture the file to a file, add the `-o` option and the filename of the file to capture to, as shown here:

```
[user@localhost ~]$ curl https://8.8.8.8 -o file.html  
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current  
                                 Dload  Upload  Total  Spent    Left  Speed  
100   216   100   216     0     0   2349      0 --:--:-- --:--:-- --:--:--  2373  
[user@localhost ~]$ cat file.html  
<HTML><HEAD><meta http-equiv="content-type" content="text/html; charset=utf-8">  
<TITLE>302 Moved</TITLE></HEAD><BODY>  
<H1>302 Moved</H1>  
The document has moved  
<A HREF="https://dns.google/">here</A>.  
</BODY></HTML>  
[user@localhost ~]$
```


Getting Help

Linux is a great alternative to the Windows operating system, but the command line can be overwhelming. However, once you learn some of these basic commands, it can be very rewarding. From time to time you may need help, and this is where the `man` command comes in handy. The `man` command, also known as the manual command, allows you to quickly look up the command arguments or usage of a command, as shown in Figure 4.25. Enter the command **man**, followed by the command you need information on, to open the manual page. For example, you can see the manual page for the `cat` command by entering **man cat** at the command line.

The manual pages are installed with any package that is installed on the system. We will only be concerned with command-line usage, but for programmers, there are manual pages that explain system calls and APIs.

Summary

This chapter provided an overview of operating systems other than Microsoft Windows. In particular, we looked at macOS and Linux, including the features and various tools included with each that appear on the CompTIA A+ 220-1202 exam.

FIGURE 4.25 The man command

```
user@server: ~
CAT (1) User Commands CAT (1)
NAME
  cat - concatenate files and print on the standard output
SYNOPSIS
  cat [OPTION]... [FILE]...
DESCRIPTION
  Concatenate FILE(s) to standard output.

  With no FILE, or when FILE is -, read standard input.

  -A, --show-all
      equivalent to -vET
  -b, --number-nonblank
      number nonempty output lines, overrides -n
  -e
      equivalent to -vE
  -E, --show-ends
      display $ at end of each line
  -n, --number
      number all output lines
  -s, --squeeze-blank
      suppress repeated empty output lines

Manual page cat(1) line 1 (press h for help or q to quit)
```

We covered the installation and uninstallation of applications in macOS. This included the various methods that can be used to download and install an application. We discussed the basic management of applications, such as creating shortcuts.

We also covered best practices for both macOS and Linux. There are best practices that technicians and administrators should follow regardless of which operating system(s) they are running, such as backup and antivirus.

The chapter concluded with some well-rounded knowledge of Linux administration. You learned a myriad of Linux commands that can be used in a variety of applications. You also learned how the Linux operating system works and how you can administer the OS. This chapter covers what you need to know for the 220-1202 exam. An entire book can be dedicated to the Linux operating system, and there is a CompTIA certification dedicated to Linux, called Linux+.

Exam Essentials

Know how applications are managed with common features and tools in macOS. Know that applications can be installed via the App Store or downloadable applications. Know

the various file types an application installation can be downloaded on. Know how applications are installed by dragging the application to the Applications folder, and be able to describe the uninstallation process of dragging the application to the Trash.

Be able to identify best practices. Best practices help keep systems and data usable. Among the best practices to follow, it is important to schedule backups, schedule disk maintenance, keep systems up to date, and make sure that antivirus/antimalware definition files are current.

Know the available System settings. You should be familiar with the various System settings. Know what each setting offers in the form of configuration for macOS.

Know how to use key macOS tools and features. A number of tools and features are included with macOS that make it an attractive operating system to use. Among them are Mission Control, Keychain, Spotlight, iCloud, Gestures, the Finder, Dock, Disk Utility, File Vault, Terminal, and Force Quit.

Know the basic Linux commands. Every version and distribution of Linux allows you to get to the command line in one way or another and interact with the shell. From there, you can give commands to navigate around (cd, pwd), to create and change file values (chmod, chown), to run commands (su, sudo), and to do many other tasks.

Know the common configuration files. The common configuration files you should know for the exam are /etc/passwd, /etc/shadow, /etc/hosts, /etc/fstab, and /etc/resolv.conf. Review the chapter for the various applications of each file.

Understand the various components of Linux. The various components of Linux that you should understand for the 220-1202 exam are the bootloader, kernel, and systemd. These components are all part of the boot process. The Linux boot process involves initializing the EFI firmware, loading the GRUB bootloader, and then loading and initializing the Linux kernel, which starts the systemd system.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Within a Linux terminal, you want to see all the files on your system in long format (using the `-l` option), including any hidden files (which requires the `-a` option). Which command should you use?
 - A. `ls -a | ls -l`
 - B. `ls -s; ls -l`
 - C. `ls -la`
 - D. `ls -a\ls -l`
2. Which of the following allows you to see all running programs in macOS?
 - A. Keychain
 - B. Mission Control
 - C. Finder
 - D. Force Quit
3. As part of your training program, you're trying to convince users to make backups on a regular basis. Which Apple app can be used to make backups of various types on a regular basis?
 - A. Time Machine
 - B. Finder
 - C. VSS
 - D. Keychain
4. Which of the following Linux commands/utilities can be used to edit a file?
 - A. `ps`
 - B. `nano`
 - C. `rm`
 - D. `ls`
5. Which of the following Linux commands/utilities can be used to edit an Ethernet connection's configuration settings?
 - A. `dd`
 - B. `apt`
 - C. `ip`
 - D. `pwd`

6. Which of the following is a macOS feature for password management in the latest version macOS 15.2 (Sequoia)?
 - A. Passwords
 - B. Keychain
 - C. Dock
 - D. Spotlight
7. The interpreter in Linux between the operating system and the user is known as the _____.
 - A. Shell
 - B. Translator
 - C. Login
 - D. GUI
8. What type of backup is kept on-site at the computer center for immediate recovery purposes?
 - A. On-path attack
 - B. Cloud copies
 - C. Journal copies
 - D. Working copies
9. Which of the following utilities can be used in Linux to download patches for installation on a workstation?
 - A. update
 - B. Shell/terminal
 - C. apt
 - D. patch
10. Which of the following commands can be used to change the owner of a file to a new owner in Linux?
 - A. cd
 - B. chmod
 - C. chown
 - D. pwd
11. Which Linux utility can be used to check and repair disks?
 - A. fsck
 - B. chkdsk
 - C. du
 - D. dumgr

12. Your iPad has an application that will not stop running. What feature/tool can you use to stop it?
- A. `kill`
 - B. Force Quit
 - C. Task Manager
 - D. Close Quit
13. Which of the following is the most common shell used with Linux?
- A. Tcl/Tk
 - B. Terminal
 - C. Bash
 - D. SSH
14. What is the name of the area at the bottom of a macOS screen where, by default, a bar of crucial icons appears?
- A. Footer
 - B. Mission Control
 - C. Taskbar
 - D. Dock
15. Which key combination can you use to bring up Spotlight from within an app?
- A. Control+Shift
 - B. Option+Tab
 - C. Command+spacebar
 - D. Alt+Home
16. Which Linux command can be used to let you run a single command as another user?
- A. `sudo`
 - B. `su`
 - C. `passwd`
 - D. `ifconfig`
17. Which of the following Linux commands will show you a list of running processes?
- A. `ls`
 - B. `cat`
 - C. `ps`
 - D. `su`

18. You are currently in a Linux terminal session and in the `/home/testuser/documents/mail` directory. Which command will take you to `/home/testuser/documents`?
- A. `cd .`
 - B. `cd ..`
 - C. `cd ...`
 - D. `cd ~`
19. If the permissions for a file are `rw-rw-r--`, what permissions apply for a user who is a member of the group to which the owner belongs?
- A. Read, write, and execute
 - B. Read and write
 - C. Read only
 - D. No access
20. What does the `-p` option with `mkdir` do?
- A. Prompts the user before creating files
 - B. Prompts the user before creating subfolders
 - C. Creates subfolders as well as folders
 - D. None of the above

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

By default, not all files and folders in a Linux directory are shown when you do an `ls` listing. Entries that start with a period (.) are considered "hidden" and not shown. Try this command in your home directory, and then compare the result with what you see when you don't use the `-a` option:

1. Type `cd /` to change to the root directory.
2. Type `ls -F` to see the files and directories in the root directory.
3. Type `ls -aF` to see everything, including hidden files.
4. Type `cd ~` to change to your home directory.
5. Type `ls -l` to see the files and directories in your home directory.
6. Type `ls -al` to see everything, including hidden files.

Chapter 5

Security Concepts

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **2.0 Security**

✓ **2.1 Summarize various security measures and their purposes.**

- Physical security
 - Bollards
 - Access control vestibule
 - Badge reader
 - Video surveillance
 - Alarm systems
 - Motion sensors
 - Door locks
 - Equipment locks
 - Security Guards
 - Fences
- Physical access security
 - Key fobs
 - Smart cards
 - Mobile digital key
 - Keys
 - Biometrics
 - Retina scanner
 - Fingerprint scanner
 - Palm print scanner
 - Facial recognition technology (FRT)
 - Voice recognition technology



- Lighting
- Magnetometers
- Logical Security
 - Principle of least privilege
 - Zero Trust model
 - Access control lists (ACLs)
 - Multifactor authentication (MFA)
 - Email
 - Hardware token
 - Authenticator application
 - Short message service (SMS)
 - Voice call
 - Time-based one-time pass word (TOTP)
 - One-time password/passcode (OTP)
- Security Assertions Markup Language (SAML)
- Single sign-on (SSO)
- Just-in-time access
 - Privileged access management (PAM)
- Mobile device management (MDM)
- Data loss prevention (DLP)
- Identity access management (IAM)
- Directory services

✓ **2.2 Given a scenario, configure and apply basic Microsoft Windows OS security settings.**

- Active Directory
 - Joining domain
 - Assigning log-in script
 - Moving objects within organizational units
 - Assigning home folders
 - Applying Group Policy



- Selecting security groups
- Configuring folder redirection

✓ **2.4 Summarize types of malware and tools/methods for detection, removal, and prevention.**

- Malware
 - Trojan
 - Rootkit
 - Virus
 - Spyware
 - Ransomware
 - Keylogger
 - Boot sector virus
 - Cryptominer
 - Stalkerware
 - Fileless
- Adware
 - Potentially unwanted program (PUP)
- Tools and methods
 - Recovery console
 - Endpoint detection and response (EDR)
 - Managed detection and response (MDR)
 - Extended detection and response (XDR)
 - Antivirus
 - Anti-malware
 - Email security gateway
 - Software firewalls
 - User education regarding common threats
 - Anti-phishing training
 - OS reinstallation



✓ 2.5 Compare and contrast common social engineering attacks, threats, and vulnerabilities.

- Social engineering
 - Phishing
 - Vishing
 - Smishing
 - QR code phishing
 - Spear phishing
 - Whaling
 - Shoulder surfing
 - Tailgating
 - Impersonation
 - Dumpster diving
- Threats
 - Denial of service (DoS)
 - Distributed denial of service (DDoS)
 - Evil twin
 - Zero-day attack
 - Spoofing
 - On-path attack
 - Brute-force attack
 - Dictionary attack
 - Insider threat
 - Structured Query Language (SQL) injection
 - Cross-site scripting (XSS)
 - Business email compromise (BEC)
 - Supply chain/pipeline attack
- Vulnerabilities
 - Non-compliant systems
 - Unpatched systems



- Unprotected systems (missing antivirus/missing firewall)
- EOL
- Bring your own device (BYOD)

✓ **2.7 Given a scenario, apply workstation security options and hardening techniques.**

- Data-at-rest encryption
- Password considerations
 - Length
 - Character types
 - Uniqueness
 - Complexity
 - Expiration
- Basic input/output system (BIOS)/Unified Extensible Firmware Interface (UEFI) passwords
- End-user best practices
 - Use screensaver locks
 - Log off when not in use
 - Secure/protect critical hardware (e.g., laptops)
 - Secure personally identifiable information (PII) and passwords
 - Use password managers
- Account management
 - Restrict user permissions
 - Restrict log-in times
 - Disable guest account
 - Use failed attempts lockout
 - Use timeout/screen lock
 - Apply account expiration dates
- Change default administrator's user account/password
- Disable AutoRun
- Disable unused services



✓ **2.9 Compare and contrast common data destruction and disposal methods.**

- Physical destruction of hard drives
 - Drilling
 - Shredding
 - Degaussing
 - Incineration
- Recycling or repurposing best practices
 - Erasing/wiping
 - Low-level formatting
 - Standard formatting
- Outsourcing concepts
 - Third-party vendor
 - Certification of destruction/recycling
- Regulatory and environmental requirements



Think of how much simpler an administrator's life was in the days before every user had to be able to access the Internet, and how much simpler it must have been when you only had to maintain a number of dumb terminals connected to a mini-tower or mainframe computer. Much of what has created headaches for an administrator since then is the inherent security risk that comes about as the network expands. As our world—and our networks—have become more connected, the need to secure data and keep it away from the eyes of those who can do harm has increased exponentially.

Realizing this, CompTIA added the Security domain to the A+ exams a number of years back. Security is now a topic that every administrator and technician must not only be aware of and concerned about, but also be actively involved in implementing methods to enforce and monitor. In the world of production, quality may be job one, but in the IT world, it is security.

This chapter, one of two chapters that focus primarily on security, will cover myriad security concepts. First, we will explore the physical aspects of security, and then we will dive deeper into the logical aspects of security. We will then look at how external forces, such as malware, social engineering, and vulnerabilities, can impact security. We will finish this chapter by looking at some common ways that you can safeguard yourself from security breaches. We will cover the proper destruction and disposal methods as well as security measures you can employ in network installations.



A+ is not the only IT certification that CompTIA offers. Security+ is one of the more popular choices. The topics found in this chapter are a subset of what you need to know for that certification.

Many of the security issues that plague networks today can be solved through the implementation of basic security elements. Some of those elements are physical (e.g., locked doors), and others are digital (e.g., antivirus/antimalware software), but all share in common the goal of keeping problems out. The following six topic areas are key:

- Physical security
- Digital security
- User education
- The principle of least privilege
- Email security
- Virtual private networks (VPNs)

As you study for the exam, know the types of physical security elements that you can add to an environment to secure it. Know, as well, what types of digital security you should implement to keep malware at bay. Understand that the first line of defense is the user. You need to educate users to understand why security is important, and you must impose the principle of least privilege to prevent them from inadvertently causing harm.

Physical Security Concepts

Physical security is the most overlooked element of security in a network. A simple lock can keep out the most curious prying eyes from a network closet or server room. A more layered approach can be implemented for higher security installations. However, the simple fact is that not a lot of time is spent on physically securing the network. In this section, we will cover the CompTIA objectives related to physical security of networks.




Exam Objective 2.1, “Summarize various security measures and their purposes,” requires you to know the various security measures for both physical and logical access.

Access Control Vestibule

The use of an access control vestibule helps to prevent nonauthorized users from tailgating. An *access control vestibule* is a small room that has two controlled doors, as shown in Figure 5.1. When a person enters the first door, they are trapped in the room until they have been authorized to enter the second controlled door. The close proximity between people in this confined space makes it uncomfortable both for the authorized user and for the nonauthorized user attempting to tailgate. In this example, the doors are controlled by radio frequency identification (RFID) readers, which we will cover later in this section.

FIGURE 5.1 A common access control vestibule setup



Badge Reader

Identification (ID) badges are used to provide proof of access. Badges can be any form of identification intended to differentiate the holder from everyone else. This can be as simple as a name badge or a photo ID. When the badge contains a photo ID, it is considered the authentication factor of something that you are.

Many ID badges also contain a magnetic strip or RFID provision so that the badge can be used in conjunction with a *badge reader*. When the information is read by the badge reader, it is sent to an access control system for authorization through the controlled door. A benefit of implementing badge readers is that it creates an electronic audit of all access to an area.

Video Surveillance

Video surveillance is the backbone of physical security. It is the only detection method that allows an investigator to identify what happened, when it happened, and, most important, who made it happen. Two types of cameras can be deployed: fixed and *pan-tilt-zoom (PTZ)*. Fixed cameras are the best choice when recording for surveillance activities. Pan-tilt-zoom (PTZ) cameras allow for 360-degree operations and zooming in on an area. PTZs are most commonly used for intervention, such as covering an area outside during an accident or medical emergency. PTZ cameras are usually deployed for the wrong reasons, mainly because they are cool! PTZs are often put into patrol mode to cover a larger area than a fixed camera can. However, when an incident occurs, they are never pointed in the area you need them! It is always best to use a fixed camera or multiple fixed cameras, unless you need a PTZ for a really good reason. They are usually more expensive and require more maintenance than fixed cameras.

Video surveillance can be deployed using two common media types: coaxial cable and Ethernet. Coaxial cable is used typically in areas where preexisting coaxial lines are in place or distances are too far for typical Ethernet. These systems are called closed-circuit television (CCTV). Coaxial camera systems generally use appliance-like devices for recording of video. These CCTV recorders generally have a finite number of ports for cameras and a finite amount of storage in the form of direct-attached storage (DAS).



Most video installations for CCTV are coaxial cable and Ethernet, as previously described. However, wireless is popular for consumer applications, such as doorbells and home surveillance cameras. These devices generally use cloud storage and require an Internet connection.

Ethernet (otherwise known as IP) surveillance is becoming the standard for new installations. Anywhere an Ethernet connection can be installed, a camera can be mounted. Power over Ethernet (PoE) allows power to be supplied to the camera, so the additional power supplies used with coaxial cameras are not needed. Ethernet also provides the flexibility of virtual local area networks (VLANs) for added security so that the camera network is isolated from operational traffic. IP surveillance uses network video recorder (NVR) software to record cameras. Because NVRs are server applications, you can use traditional storage such as network area storage (NAS) or storage area network (SAN) storage. This allows you to store and treat the video recordings like traditional data.

Coaxial camera networks can be converted to IP surveillance networks with the use of a device called a *media converter*. These devices look similar to a CCTV recorder. They have a limited number of ports for the coaxial cameras and are generally smaller than the CCTV recorder. This is because they do not have any DAS. The sole purpose of the media converter is to convert the coaxial camera to an Ethernet feed to the NVR.

The use of IP video surveillance allows for a number of higher-end features such as camera-based motion detection, license plate recognition (LPR), and motion fencing. Advanced NVR software allows cameras to send video only when motion is detected at the camera; this saves on storage for periods of nonactivity. LPR is a method of detecting and capturing license plates in which the software converts the plate to a searchable attribute for the event. With motion fencing, an electronic fence can be drawn on the image so that any activity within this region will trigger an alert. Among the many other features are facial recognition and object recognition.

Motion Sensors

There are several different motion sensor types that you can use to detect unauthorized access. Passive infrared (PIR) is the most common motion detection used today, mainly because of price. PIR sensors operate by monitoring the measurement of infrared radiation from several zones. In Figure 5.2, you can see the reflective panel that divides the infrared zones. A PIR sensor will always have this grid pattern on the sensor's face.

FIGURE 5.2 A typical PIR sensor



Microwave detectors also look like PIR sensors, but they do not have a reflective panel. Microwave detectors are common in areas where wide coverage is needed. They operate by sending pulses of microwaves out and measuring the microwaves received. These detectors are more expensive than PIR sensors and are susceptible to external interference, but they have a wider area of coverage.

Vibration sensors are another type of sensor used for motion detection. Although you may have seen them in the latest over-the-top heist movie, vibration sensors are really used in physical security systems. They are most often implemented as seismic sensors. They help protect from natural disasters and accidental drilling, or the occasional over-the-top heist.

Alarm System

An alarm system is another type of physical security system. It provides a method to alert security personnel in the event of unauthorized access or a break in. An alarm system can be configured to trigger in the event of an access control system logging unauthorized access to a controlled door. However, it is more common to find alarm systems installed for break-in detection and response.

An alarm system can be configured to use motion sensors, video surveillance, magnetic contacts, and a multitude of other sensors. Each sensor will be installed in a different logical zone. The perimeter of the building might be zone 1, the server room might be zone 2, and so on. There could be many different zones configured for each sensor. The main purpose of the zone is to communicate the location of the sensor being tripped so that law enforcement agents can respond to that location.

A monitoring company is typically contracted with the installation. The purpose of the monitoring company is to act as a buffer between the zone being tripped and a law enforcement agency. The alarm panel will dial out to a monitoring station and it will transmit the account number and the zone that is tripped. The monitoring station can then process its call-down list. The call-down list will typically consist of the phone number of the supervisor of the area, the in-house security personnel, and ultimately a law enforcement agency. The monitoring company can also monitor the health of the alarm panel, depending on the model of alarm panel and its features.

Door Locks

The most common physical prevention tactic is the use of locks on doors and equipment. This might mean the installation of a tumbler-style lock or an elaborate electronic combination lock for the switching closet. If a tumbler-style lock is installed, then the appropriate authorized individuals who require access will need a physical key. Using physical keys can become a problem, because you may not have the key with you when you need it the most, or you can lose the key. The key can also be copied and used by unauthorized individuals. Combination locks, also called cipher locks, can be reprogrammed and do not require physical keys, as shown in Figure 5.3. Combination locks for doors can be purchased as mechanical or electronic.

FIGURE 5.3 A typical combination door lock

When physical locks use keys, the factor of authentication is considered something that you have—because you must have the key. When physical locks use ciphers, the authentication is considered something you know—because you must know the cipher.

Equipment Locks

There are many different types of *equipment locks* that can secure the information and the device that holds the information. Simply thwarting the theft of equipment containing data and restricting the use of USB thumb drives can secure information. In the following sections, we will cover several topics that are directly related to the physical aspects of information security.

Cable Locks

Cable locks are used to secure laptops and any device with a Universal Security Slot (USS), as shown in Figure 5.4. A cable lock is just that—a cable with a lock at one end. The lock can be a tumbler or a combination, as shown in Figure 5.5. The basic principle is that the end of the lock fits into the USS. When the cable is locked, the protruding slot of metal turns into a cross that cannot be removed. This provides security to expensive equipment that can be stolen due to its portability or size.

FIGURE 5.4 A Universal Security Slot**FIGURE 5.5** A standard cable lock

Server Locks

Most servers come with a latch-style lock that prevents someone from opening the server, but the tumbler-style lock is trivial to open. Anyone with a paperclip can open these locks if they have forgotten the keys. Other types of server locks are holes for padlocks that latch through the top cover and the body of the server. However, over the past 10 years, a

declining number of servers come with this feature. This is mainly due to the fact that servers can be better secured behind a locked rack-mounted enclosure. Rack-mounted enclosures generally come with a tumbler-style lock that can protect all the servers and network equipment installed in the cabinet, while still providing airflow.

USB Locks

Universal Serial Bus (USB) locks can be put into place to physically lock out USB ports on a workstation or server from use. These devices are extremely rare to find, because most equipment and operating systems allow for the USB ports to be deactivated. USB locks work by inserting a small plastic spacer into the USB port. Once inserted, the spacer latches to the USB detent with plastic teeth. A tool is required to remove the USB spacer.

Security Guards

Physical security begins with personnel—specifically, security-focused personnel, such as security guards. *Security guards* should be responsible for limiting access from the outer perimeter of your installation. Security guards typically use photo IDs, also known as ID badges, to allow access to the installation. Exceptions to this are people on the entry control roster; in some secured buildings, only people on the entry control roster are allowed to enter. In this type of scenario, the ID badge is used only to provide ID. This is common in government and sensitive installations.

Fences

Fences are a physical security barrier to keep unauthorized persons out of a secure area. Exterior fences can be arranged so that they create a choke point where a guard can inspect credentials to allow authorized personnel into the area. Guards can also be replaced with electronic locks and RFID readers to limit access. When installed in conjunction with a video camera system to surveil the area around the entry point, a fence creates a very secure outer layer for your facility. Fences should be considered the outermost security layer of a multiple-barrier system.

In addition to an exterior fence, the building that houses the data should have RFID readers and electronic door locks. The innermost area surrounding the equipment can also be segmented with a fence and additional access controls, such as standard keyed locks or electronic access control. When fences are used in the interior of the data center, air quality can be maintained while preventing unauthorized access, as shown in Figure 5.6.

Using multiple barriers as described allows contractors for HVAC systems to maintain their systems, while preventing direct physical access to the servers and equipment.

FIGURE 5.6 Interior data center fences



Bollards

A *bollard* is an architectural structure that acts as a visual indicator for a perimeter. They are also very sturdy, since their second function is to act as a barrier for the perimeter and protect the area. They are commonly found around areas where a truck or other vehicle can cause damage. In Figure 5.7, the bollard is protecting a fiber-optic vault from accidental damage by a vehicle. Bollards can also be found in the interior of a building if there is potential for damage to a protected area from a vehicle, such as a forklift or equipment cart.

FIGURE 5.7 A typical bollard

Physical Security for Staff

Organizations should authorize and audit staff access for sensitive areas inside a facility. Implementing physical security for staff is one way you can control access to the physical equipment and the data that is stored on the equipment. By limiting access to the equipment and the underlying data, you can prevent service disruptions or loss of data. This section will focus on ways that you can implement physical security for staff.

Key Fobs

Key fobs are named after the chains that used to hold pocket watches to clothes. *Key fobs* are embedded *radio frequency identification (RFID)* circuits that fit on a set of keys and are used with physical access control systems, as shown in Figure 5.8. They are often used for access to external and internal doors for buildings. Key fobs are close-proximity devices that authorize the user for entry; an electronic lock is actuated when the device is presented, and the door can be opened. This is an authentication factor of something that you have.

FIGURE 5.8 A key fob

Smartcards and RFID Badges

A *smartcard* is the size of a credit card with an integrated circuit embedded into the card (also called an integrated circuit chip [ICC]). The chip is exposed on the face of the card with surface contacts, as shown in Figure 5.9. Smartcards are used for physical authentication to electronic systems and access control systems and require a PIN or password. A smartcard is considered a multifactor authentication method because it is something you have (card) and something you know (PIN or password). The U.S. military uses smartcards called Common Access Cards (CACs) for access to computer systems and physical access controls.

FIGURE 5.9 A typical smartcard



An RFID badge is a wireless, no-contact technology used with RFID transponders. RFID badges typically work on the 125 kHz or 13.56 MHz radio frequency and are passively powered by the RFID transponder. When an RFID badge is placed in close proximity to the RFID transponder, the radio frequency (RF) energy emitted by the transponder powers a chip in the RFID badge. The RFID chip then varies the frequency back to the transponder in the effort to transmit its electronic signature (number). This type of authentication is considered something you have.

Keys

Physical keys are extremely hard to control and do not allow for the auditing of their usage. A physical key can be lent to someone, copied, stolen, or used by an unauthorized person. Because of the problems surrounding physical keys, their use should largely be avoided.

If keys are absolutely necessary, then a two-person system should be considered. A two-person system requires that two people must use their keys to open one lock, although nothing stops the keys from being lent to the same person to open a door.

Another option is to use an electronic lock box for management of the keys. When a technician needs a particular key, they will log into the key box and check out the key needed. This system allows for auditing controls, but it does not prevent copying of keys.

Mobile Digital Key

The history of mobile digital keys displays the rapid technological advancements and evolving security needs, beginning in the early 2010s with the hospitality industry. Inter-Continental Hotels Group (IHG) first experimented with mobile keys in 2010, followed by Starwood Hotels' SPG Keyless in 2014, which allowed guests to unlock rooms with smartphones. Major hotel chains like Hilton and Marriott quickly adopted and refined this technology, enhancing guest experiences and setting new security standards. By 2019, Hilton had implemented digital keys in over 4,000 hotels, while Marriott developed a user-friendly system leveraging Starwood's legacy technology.

Today, the use of mobile digital keys extends far beyond hotels. They are now commonly used in residential buildings, offices, and even vehicles. Digital keys utilize technologies like Bluetooth, NFC (near-field communication), and Wi-Fi to communicate securely with locks. This wireless communication allows users to unlock doors with their smartphones, providing a seamless and convenient access experience.

Mobile device keys offer superior convenience and security compared to traditional keys. They eliminate the need to carry physical keys, reduce the risk of lockouts, and provide seamless access through smartphones. Enhanced security features like encryption and two-factor authentication protect against unauthorized access, while the ability to remotely deactivate keys if a device is lost or stolen adds an extra layer of protection. This combination of convenience and robust security makes mobile device keys a better choice for modern access control.

Biometrics

Biometrics refers to the technology that identifies individuals based on their unique physical or behavioral traits. This method of identification is increasingly popular due to its high level of security and convenience. This factor of authentication is considered "something that you are," and we will discuss other factors later in this chapter.

In more recent years, most mobile phones have implemented biometrics in the access control of the mobile device. Several manufacturers have adopted fingerprint access control, and some have even adopted facial recognition via the forward-pointing camera.

Biometric systems typically follow three main steps: enrollment, where the system captures and stores a biometric sample; storage, where the sample is converted into a digital format and stored in a database; and comparison, where the system compares a new sample with the stored data to verify identity. These systems are designed to be difficult to forge or

steal, making them a reliable form of authentication. The following are the various types of biometric systems outlined for the 220-1202 exam:

Retina Scanner Retina scanners work by using low-energy infrared light to scan the unique pattern of blood vessels in the retina, located at the back of the eye. The light traces a path on the retina, and the differences in light absorption by the blood vessels create a unique pattern that is digitized and stored. Retinal patterns are highly stable and unique, making them difficult to replicate. To safeguard against false positives, retina scanners use sophisticated algorithms to ensure the scanned pattern matches the stored data with high precision.

Fingerprint Scanner Fingerprint scanners capture the unique patterns of ridges and valleys on a person's fingertip. There are several types of fingerprint scanners, including optical, capacitive, and ultrasonic. Optical scanners use light to capture an image of the fingerprint, while capacitive scanners use electrical currents to map the ridges and valleys. Ultrasonic scanners use sound waves to create a detailed image. These scanners compare the captured fingerprint with stored data to verify identity. Figure 5.10 shows a typical biometric device.

FIGURE 5.10 A fingerprint biometric lock



Palm Print Scanner Palm print scanners work similarly to fingerprint scanners but capture the larger and more complex patterns found on the palm. These patterns include ridges, lines, and creases that are unique to each individual. Palm print scanners use

imaging sensors to capture these details and convert them into a digital format for storage and comparison.

Facial Recognition Technology (FRT) Facial recognition technology identifies individuals by analyzing the unique features of their faces. This includes the distance between the eyes, the shape of the cheekbones, and the contour of the lips, among other features. The system captures an image of the face, converts it into a digital format, and compares it with stored data.

Voice Recognition Technology Voice recognition technology identifies individuals based on their unique vocal characteristics, such as pitch, tone, and cadence. The system captures a sample of the person's voice, converts it into a digital format, and analyzes it for unique patterns. To safeguard against false positives, voice recognition systems use advanced algorithms that can distinguish between similar-sounding voices and account for variations in speech due to illness or emotional state.

With the passing of time, the definition of biometrics is expanding from simply identifying physical attributes about a person to being able to describe patterns in their behavior. Recent advances have been made in the ability to authenticate someone based on the key pattern that they use when entering their password (how long they pause between keys, the amount of time each key is held down, and so forth). An organization adopting biometric technologies needs to consider the controversy they may face. Some authentication methods are considered more intrusive than others. The error rate also needs to be considered, along with an acceptance of the fact that errors can include both false positives, where the reader allows access falsely, and false negatives, where the reader denies access erroneously. Therefore, biometrics are often used with another factor of authentication, such as a PIN number. This approach provides multifactor authentication.

Lighting

Most security cameras work on the principle of collecting light to record a picture. As light levels decrease, the quality of the picture decreases significantly. Therefore, areas in which you have video cameras should have sufficient levels of lighting. In reality, any area that is sensitive should have a level of lighting, since threat agents often hide in darker areas to avoid being seen.

Lighting the area doesn't always require a visible light source. Most cameras sensors built in the last 10 years allow for the collection of light from infrared (IR) light sources. This means that even in the dark you can surveil and record an area, although a visible light source does deter unauthorized access.

Magnetometers

The *magnetometer*, also known as a metal detector, uses an electromagnetic field to detect metallic objects. We have all seen these devices at a choke point in the airport or government building. When a metal detector is used for people entering a facility, you can detect

weapons, such as guns or knives. When a metal detector is deployed in this fashion, it protects your staff from threat agents with malicious intent.

A metal detector can also be used to monitor people leaving a facility. The metal detector can monitor staff leaving with equipment. When it is used in this way, it will protect against data loss and theft. However, unless your organization is regulated as a highly classified facility, this approach will be hard to enforce and could infringe on your employees' privacy.

Logical Security

Whereas the topic of physical security concepts, from CompTIA's standpoint, focuses on keeping individuals out, logical security focuses on keeping harmful data and malware out as well as on authorization and permissions. This logical security includes devices and methods that protect the environment logically, such as firewalls, antivirus software, and directory permissions, just to name a few. The areas of focus are antivirus software, firewalls, antimalware, user authentication/strong passwords, and directory permissions. Each of these topics is addressed in this section.

Principle of Least Privilege

The *principle of least privilege* is a common security concept that states a user should be restricted to the fewest number of privileges that they need to do their job. By leveraging the principle of least privilege, you can limit internal and external threats. For example, if a frontline worker has administrative access on their computer, they have the ability to circumvent security; this is an example of an internal threat. Along the same lines, if a worker has administrative access on their computer and receives a malicious email, a bad actor could now have administrative access to the computer; this is an example of an external threat. Therefore, only the permissions required to perform their tasks should be granted to users, thus providing least privilege.

Security is not the only benefit to following the principle of least privilege, although it does reduce the surface area of attack because users have less access to sensitive data that can be leaked. When you limit workers to the least privilege they need on their computers or the network, fewer intentional or accidental misconfigurations will happen that can lead to downtime or help desk calls. Some regulatory standards require following the principle of least privilege. By following the principle of least privilege, an organization can improve on compliance audits by regulatory bodies.

Zero Trust

In a traditional trust model, everything within the perimeter of the network is trusted and everything outside is not. When an external user needs to access an internal resource, they connect to the internal network via a virtual private network (VPN). Connecting to the

network with a VPN makes the device they are connecting with trusted because the user is authenticated with their credentials.

The traditional model works well when the expectation is that you drive to work and log on to a trusted computer to access sensitive information. However, with today's mixture of work from home and mobile device access, the traditional model doesn't always work. The traditional model also doesn't address cloud resource access. This is where the Zero Trust model excels. With a Zero Trust model, nothing is trusted regardless of the location and everything is authenticated. If a user logs on to a cloud resource, the device is authenticated and the user is authenticated before they can connect to the data or application. This prevents a bad actor from using a user's compromised device to access sensitive data. Every vendor has a different implementation of Zero Trust, but the concept remains the same; nothing is trusted and everything and everyone must be authenticated.

Access Control Lists

Access control lists (ACLs) are used to control traffic and applications on a network. Every network vendor supports a type of ACL method; for the remainder of this section, I will focus on Cisco ACLs.

An ACL method consists of multiple access control entries (ACEs) that are condition actions. Each entry is used to specify the traffic to be controlled. Every vendor will have a different type of control logic. However, understanding the control logic of the ACL system allows you to apply it to any vendor and be able to effectively configure an ACL. The control logic is defined with these simple questions:

- How are the conditions of an ACL evaluated?
- What is the default action if a condition is not met?
- How is the ACL applied to traffic?
- How are conditions edited for an ACL?

Let's explore the control logic for a typical Cisco layer 3 switch or router. The conditions of the ACL are evaluated from top to bottom. If a specific condition is not met for the ACL, the default action is to deny the traffic. Only one ACL can be configured per interface, per protocol, and per direction. When you are editing a traditional standard or extended ACL, the entire ACL must be negated and reentered with the new entry. With traditional ACLs, there is no way to edit a specific ACL on the fly. When editing a named access list, each condition is given a line number that can be referenced so that the specific entry can be edited. For the remainder of this section, we will use named access lists to illustrate an applied access list for controlling traffic.


In Figure 5.11 you can see a typical corporate network. There are two different types of workers: HR workers and generic workers. We want to protect the HR web server from access by generic workers.

We can protect the HR server by applying an ACL to outgoing traffic for Eth 0/0 and describing the source traffic and destination to be denied. We can also apply an ACL to the incoming interface of Eth 0/2 describing the destination traffic to be denied. For this

example, we will build an access list for incoming traffic to Eth 0/2, blocking the destination of the HR server.

```
Router(config)# ip access-list extended block-hrserver
Router(config-ext-nacl)# deny ip any host 192.168.1.4
Router(config-ext-nacl)# permit ip any any
Router(config-ext-nacl)# exit
Router(config)# interface ethernet 0/2
Router(config-if)# ip access-group block-hrserver in
```

FIGURE 5.11 A typical corporate network



This ACL, called `block-hrserver`, contains two condition action statements. The first denies any source address to the specific destination address of 192.168.1.4. The second allows any source address to any destination address. We then enter the interface of Eth 0/2 and apply the ACL to the inbound direction of the router interface. The rule will protect the HR server from generic worker access while allowing the generic workers to access all other resources and the Internet.

It is important to note that the focus of this section is to understand how ACLs are used to protect resources. It is not important to understand how to build specific ACLs, since commands will be different from vendor system to vendor system.



Access control lists can also describe the logical controls for file access. These types of ACLs control who has access with which files. We will cover file-based ACLs in Volume 2 Chapter 6, “Securing Operating Systems.”

Authentication Factors

This section discusses various components of physical security that control access. When components are used to control access, they do so based on the authentication of people. Authentication can happen in several different ways:

Something You Know This is commonly a username and password or *personal identification number (PIN)*. You can make passwords more complex by requiring uppercase, lowercase, numeric, and symbol combinations.

Something You Have This is based on something you physically have. This chapter covers several types of devices, including key fobs, keys, and smartcards, just to name a few.

Something You Are This authentication factor is based on something that makes you physically unique. Biometrics such as fingerprints, palm, face, your voice, and retina scans should be the first things that come to mind.

Somewhere You Are A relatively newer factor of authentication is based on somewhere you are. With the proliferation of global positioning system (GPS) chips, your current location can authenticate you for a system. This is performed by authentication rules based on the user’s location.

Something You Do Another relatively newer factor of authentication for network systems is based on something you do. Although it has been used for hundreds of years for documents and contracts, a signature is something that you do and don’t even think about how you do it. It is unique to you and only you, because there is a specific way you sign your name. Typing your name into the computer is something you do and don’t think about, but there is a slight hesitation that you make without knowing it.

Multifactor Authentication

All authentication is based on something that you know, have, are, or do, or a location you are in. A common factor of authentication is a password, but passwords can be guessed, stolen, or cracked. A fingerprint can be lifted with tape, a key can be stolen, or a location spoofed. No one factor is secure by itself because it can be compromised easily.

When more than one item (factor) is used to authenticate a user, this is known as *multi-factor authentication* (MFA). It may take two, three, or four factors to authenticate, but as long as it is more than one, as the name implies, it is known as multifactor. One of the most common examples where this is used in everyday life is at an ATM. In order to withdraw money, a user must provide a card (one factor) and a PIN number (a second factor). If you know the PIN number but do not have the card, you cannot get money from the machine. If you have the card but do not have the PIN number, you cannot get money from the machine.

In this section we will cover the most common two-factor (2FA)/multifactor authentication methods used by protected applications. The following methods are generally used in conjunction with a traditional user and password combination. It should be assumed that when we talk about 2FA it provides the same functionality as MFA.

Email

Some applications use email as a 2FA method. However, using email as a 2FA option is probably the least secure method. This is mainly due to the fact that people reuse passwords. If your banking website username and password is compromised (something you know) and you reuse the same credentials on email, it provides no protection. Hopefully the email account is protected with 2FA in a way that it requires something you have.

Email is useful as a notification method when someone logs into a secure login. However, keep in mind the threat agents know this as well. If your email account is compromised, a threat agent will often create a rule in your email box to dump these notifications directly to the trash.

Short Message Service (SMS)

Some applications will allow the use of short message service (SMS) text messages as the 2FA method. When this method is used, a simple text message is sent to the user's phone number. The message will contain a random 5- to 8-digit code that the user will use to satisfy the 2FA requirement. When you first set up this 2FA method, the protected application will request the code before turning on 2FA. This is done to verify that the phone number is correct and that you can receive text messages.

Voice Call

Some applications that are protected by 2FA will allow voice calls to be initiated to the end user. This is usually done if the person does not have a phone that accepts text messages. The voice call will recite a 5- to 8-digit code that the user will use to satisfy the 2FA requirement. This process is similar to SMS, with the difference being it is an automated voice call.

Hardware and Software Tokens

A hard token is a physical device, often resembling a key fob or a USB thumb drive. All of these types of devices serve as a “something you have” authentication factor, providing an additional layer of security. SecurID from RSA is one of the best-known examples of a physical hardware token in the form of a key fob, as shown in Figure 5.12.

FIGURE 5.12 An RSA security key fob

The various hardware tokens operate differently based upon the type of hardware token being used and the implementation. The following are the various token systems you might encounter.

Time-based One-Time Password (TOTP) Time-based One-Time Password (TOTP) is a dynamic authentication method that generates a unique, time-sensitive code based on a shared secret key and the current time. When a user attempts to log in, the TOTP algorithm combines the secret key with the current timestamp to produce a one-time password that is valid for a short period, typically 30 to 60 seconds. This code is then entered by the user to complete the authentication process. The time-sensitive nature of TOTP codes enhances security by ensuring that even if a code is intercepted, it quickly becomes invalid. This method is widely used in two-factor authentication (2FA) systems, providing an additional layer of security by requiring both something the user knows (a password) and something the user has (the TOTP code).

One-Time Password/Passcode (OTP) One-Time Password (OTP) is a static code that is generated for a single login session or transaction. Unlike TOTP codes, OTPs are not time-dependent and can be delivered through various channels such as SMS, email, or hardware tokens. When a user initiates a login or transaction, an OTP is generated and sent to the user, who then enters the code to complete the process. Because OTPs are not time-sensitive, they can be vulnerable if intercepted and used before they expire. Despite this, OTPs remain a popular choice for enhancing security in various applications, particularly where time-based constraints are less critical.

USB Tokens USB tokens, such as Fast Identity Online (FIDO) keys, are another popular form of hardware token. These devices plug directly into a computer's USB port and use cryptographic techniques to authenticate the user. When a user attempts to log in, they simply insert the USB token and press a button on the device to complete the authentication process. This method is highly secure as it requires physical possession of the token and cannot be easily replicated.

A new type of hardware token is becoming the new standard, and it is really considered a software token or soft token. It operates the same as a hardware token, but it is an application on your cell phone that provides the code. Google Authenticator is one example of

these types of applications. Microsoft also has an authenticator application called Microsoft Authenticator that is similar to Google Authenticator.

When configuring 2FA on an application, you have two ways of adding an account to the authenticator application. You can take a picture of a quick response (QR) code, or you can enter a security code into the authenticator application. If you choose to use a QR code, then the application turning the 2FA on will present a QR code that can be scanned by the authenticator application. If you choose to use a setup key, the application turning on the 2FA will provide a key. There is generally a second step before the application is protected by 2FA, where you will be required to enter a code from the authenticator application to the protected application. A lengthy one-time-use backup key is also generated, in case you need to turn 2FA off because your device is lost or stolen.

Single Sign-On (SSO)

Single sign-on (SSO) is an authentication method that allows users to access multiple applications and services with a single set of login credentials. This approach simplifies the user experience by eliminating the need to remember and manage multiple passwords, reducing the likelihood of password fatigue and the associated security risks. SSO is becoming increasingly popular due to its ability to streamline access management, enhance security, and improve productivity. By centralizing authentication, organizations can enforce stronger security policies and reduce the burden on IT support teams, as users are less likely to forget their credentials or require password resets.

SSO operates based on a trust relationship between an identity provider (IdP) and the service provider (SP). The IdP might be Windows Active Directory or Microsoft Entra ID which are two popular IdPs today. The SP is typically the cloud-based application that requires the authentication. When a user attempts to access the application, the service provider redirects the user to the identity provider for authentication. The identity provider verifies the user's credentials and, upon successful authentication, generates a token that is sent back to the service provider. This token contains information about the user's identity and is used to grant access to the requested application. This process allows users to log in once and gain access to multiple applications without needing to re-enter their credentials. The token is often stored in the user's web browser in the form of a web browser cookie.

When an account in the IdP is deactivated, access to all associated applications and services are also deactivated. This centralized control ensures that once a user's account is disabled, they can no longer access any of the linked applications, thereby enhancing security and preventing unauthorized access. This feature is particularly useful in scenarios such as employee offboarding, where it is crucial to promptly revoke access to sensitive information and systems.



Some applications will constantly check the validity of the token, and other applications will just check to see if the web browser cookie is still valid. It is always advisable to revoke the token during an employee offboarding so the token is invalidated, even if it still lives in someone's web browser cache.

Security Assertion Markup Language (SAML)

Security Assertion Markup Language (SAML) is a standard protocol used to facilitate SSO by enabling the exchange of authentication and authorization data between the identity provider and service providers. SAML works by passing XML-based authentication information between the IdP and SPs. When a user tries to access a service, the service provider sends a SAML request to the identity provider. The IdP authenticates the user and responds with a SAML assertion, which includes the user's identity and authorization information. The service provider then uses this assertion to grant the user access to the application. SAML enhances security by ensuring that authentication data is transmitted securely and that users need to authenticate only once to access multiple services.

Conditional Access

Conditional access is a security framework that enhances the protection provided by SSO by enforcing specific access controls based on predefined conditions. These conditions can include factors such as user location, device compliance, and real-time risk detection. By evaluating the context of access requests, conditional access ensures that only authenticated and authorized users can access resources. For example, an organization might require MFA for users attempting to log in from unfamiliar locations or block access from devices that are noncompliant with security policies. This adaptive approach to security helps mitigate risks associated with lost or stolen devices and unauthorized access attempts, further strengthening the overall security posture.

Just-in-Time Access

Just-in-time access is a security strategy designed to minimize the risk of unauthorized access by granting users elevated privileges only for the duration necessary to complete a specific task. This approach aligns with the principle of least privilege, ensuring that users have the minimum level of access required to perform their duties. By reducing the window of time during which elevated privileges are active, just-in-time access helps to mitigate potential security threats and prevent misuse of sensitive information.

Microsoft Entra ID Governance plays a crucial role in supporting just-in-time access by automating and streamlining identity and access management processes. With Entra ID Governance, organizations can define and enforce policies that ensure the right individuals have the appropriate access at the right times. This includes setting up approval workflows, access reviews, and automated expiration of access rights. As an example, imagine an employee needs temporary access to a sensitive financial system to complete a quarterly report. They submit a request through the Entra ID Governance portal, detailing the reason and duration needed. The request goes to their supervisor for approval. Their supervisor reviews and approves it, granting the employee access for 48 hours. After 48 hours, the employee's access is automatically revoked. The supervisor or an administrator can later review access logs to ensure compliance. This process ensures the employee gets the right access at the right time while maintaining security and compliance. The approval process is not bound only to a supervisor's approval; it can be automated based on a number of conditions.

Privileged Access Management (PAM)

Privileged access management (PAM) is a comprehensive security strategy designed to control, monitor, and protect privileged accounts within an organization. These accounts have elevated permissions that allow users to perform critical administrative tasks, access sensitive information, and make system changes. PAM solutions work by enforcing strict controls over who can access these accounts, monitoring their use, and ensuring that only authorized individuals can perform high-level administrative tasks. This approach helps to mitigate the risks of credential theft and privilege misuse, which can lead to significant security breaches.

Microsoft Entra ID Privileged Identity Management (PIM) is a robust tool that enhances PAM by providing granular control over privileged access within Microsoft environments. PIM allows organizations to manage, control, and monitor access to critical resources in Microsoft Entra ID, Azure, and other Microsoft Online Services such as Microsoft 365 and Microsoft Intune. Key features of PIM include just-in-time access, which grants temporary permissions to users only when needed, and time-bound access, which limits the duration of access to sensitive resources. Additionally, PIM enforces multifactor authentication for role activation, ensuring that only verified users can gain elevated access.

Mobile Device Management (MDM)

The traditional workforce has become a mobile workforce, with employees working from home, on the go, and in the office. Mobile devices such as laptops, tablets, and smartphones are used by employees to connect to the organization's cloud resources. *Bring your own device (BYOD)* has been embraced as a strategy by organizations to alleviate the capital expense of equipment by allowing employees to use devices they already own.

Because employees are supplying their own devices, a formal document called the BYOD policy should be drafted. The BYOD policy defines a set of minimum requirements for the devices, such as size and type, operating system, connectivity, antivirus/antimalware solutions, patches, and many other requirements the organization will deem necessary.

Many organizations use *mobile device management (MDM)* software that dictates the requirements for the BYOD policy. MDM software helps organizations protect their data on devices that are personally owned by the employees. When employees are terminated or a device is lost, the MDM software allows a secure remote wipe of the company's data on the device. The MDM software can also set policies requiring passwords on the device. All of these requirements should be defined in the organization's BYOD policy.

Data Loss Prevention

When an organization suffers data loss, the information is not lost—it is just out of the control of the organization. Data loss prevention (DLP) is a system used to ensure that end users do not send sensitive and/or critical information outside the organization. An example of sensitive information is a Social Security number or the secret sauce recipe for a hamburger chain. The user may not even know they are doing anything wrong in many instances.

DLP systems ensure that data loss will not occur because the damaging effects can be long term. A common DLP system is an outbound email filter. The email transport filters can sanitize or restrict information such as credit card numbers, Social Security numbers, or any type of data they are programmed to detect. Microsoft Exchange Server allows administrators to create custom DLP policies or use prebuilt DLP policies that ship with Microsoft Exchange. Other systems can classify data based on their sensitivity and restrict users from emailing the information to users outside the company or uploading the information to public cloud services.

DLP systems are efficient at identifying potential data loss and they provide a hard control for administrators, but it is best practice to implement soft controls as well by working with the organization's HR department to develop policies and effective training for end users.

Identity and Access Management

Identity and access management (IAM) is a security framework used for the authentication and authorization of users. The IAM model has been adopted by service providers as a means of managing users with their application. The IAM framework allows a service provider to define the authentication of users by means of a username and password, certificate keypair, or even single sign-on (SSO), just to name a few methods.

The framework also defines the authorization of user access for a resource within the service provider. Many service providers have adopted a model of security in which the object and accompanying permission is defined. This allows a granular control of what can be accessed by the user and what can be done with the service or object by the user. For example, if a database is created on a service provider, the specific database can be secured so that a particular user has access only to read objects within the database. However, administrators might have a higher level of permissions to alter information.

Directory Services

Logical security and directory services work together to enhance the overall security and management of network resources. Logical security protects data and resources through software-based measures, such as authentication, authorization, encryption, and access controls. Directory services, on the other hand, provide a centralized repository for storing and managing information about network resources, including users, computers, and services.

Authentication and Authorization Directory services like Microsoft Active Directory (AD) and Lightweight Directory Access Protocol (LDAP) play a crucial role in authentication and authorization processes. When a user attempts to access a network resource, the directory service verifies their identity (authentication) and checks their permissions (authorization) to ensure they have the appropriate access rights. This helps prevent unauthorized access and ensures that only legitimate users can access sensitive information.

Access Control Directory services enable administrators to define and enforce access control policies. By organizing users into groups and assigning permissions based on roles, directory services streamline the management of access rights. This ensures that users have access only to the resources they need, reducing the risk of data breaches and unauthorized access.

Policy Enforcement Directory services allow administrators to implement and enforce security policies across the network. For example, in an Active Directory environment, Group Policy Objects (GPOs) can be used to enforce password policies, software installation rules, and other security settings. This centralized management simplifies the enforcement of security policies and ensures consistency across the network.

Audit and Monitoring Directory services often include auditing and monitoring capabilities that help track user activities and access patterns. By logging authentication attempts, access to resources, and changes to directory information, administrators can detect and respond to suspicious activities, enhancing the overall security posture of the network.

By integrating logical security measures with directory services, organizations can create a robust security framework that protects their network resources, ensures compliance with security policies, and provides a seamless user experience.



Exam Objective 2.2, “Given a scenario, configure and apply basic Microsoft Windows OS security settings,” requires you to know the basics of administration for Active Directory (AD). You should understand the various Active Directory components and their features.

Active Directory

Microsoft originally released *Active Directory (AD)* with Windows 2000 Server to compete with Novell Directory Services (NDS). Active Directory is a highly scalable directory service that can contain many different objects, including users, computers, and printers, just to name a few. Active Directory uses the LDAP protocol to quickly look up objects. It’s important to understand that Active Directory is *not* the authentication mechanism; it is only the directory for storing and for the lookup of objects. Active Directory works in conjunction with Kerberos, which is the protocol that performed the authentication of users.

Active Directory uses a directory partition called the *schema partition* to describe classes of objects and the attributes that define each object. Each attribute is defined in the schema based on the value it will contain. For example, a user class can have a first name, last name, middle initial, description, and a number of other attributes. A specific user account is then created and configured in the GUI Microsoft Management Console (MMC) called Active Directory Users and Computers. There are also several attributes that don’t show up in the


management tool, such as when an object was last changed or replicated, as well as many other attributes used by Active Directory for management.

Domain

A domain is a hierarchical collection of security objects, such as users, computers, and policies, among other components. Active Directory domains are named with a Domain Name System (DNS) name. For example, `sybex.com` would be the root domain; if you wanted to add a new domain, you would append the namespace to the left, as follows: `east.sybex.com`. Using a DNS namespace is one of the ways that Active Directory is scalable and hierarchical. Many organizations never need anything more than one domain to contain all their security objects.

When a user authenticates against an Active Directory domain, a domain access token is issued, as shown in Figure 5.13. You can consider these to be keys for the various locks (ACLs) on resources. If the user is a member of a particular security group, it will be in their security token. When the user encounters a file that is secured with an ACL, the security token is presented. If there is a matching credential, then the user is granted the associated file permission on the ACL.

FIGURE 5.13 Active Directory security tokens



Joining Active Directory Domain

Joining an Active Directory (AD) domain offers numerous benefits, particularly for organizations seeking centralized management and enhanced security. When a computer joins an AD domain, it becomes part of a network where administrators can manage user accounts, enforce security policies, and deploy software updates from a single point of control. This centralized approach simplifies IT management, reduces administrative overhead, and ensures consistent policy enforcement across all devices. Additionally, users benefit from single sign-on (SSO) capabilities, allowing them to access various network resources with one set of credentials.

When a computer joins an AD domain, it establishes a trust relationship with the domain controller, which authenticates the computer and creates a computer account in the AD database. This trust allows the computer to securely communicate with the domain controller and other domain resources. When a domain user logs on, the computer forwards their

credentials to the domain controller for verification. If valid, the domain controller generates an access token, enabling the user to access network resources without re-entering credentials.



To learn how to join Windows 11 to an Active Directory domain, refer to the “Networking Models” section in Volume 2 Chapter 3, “Windows Administration.” You will find step-by-step instructions on joining a domain in this section.

Organizational Units

A domain can hold security objects, but you need to have some organization to the many different objects that you will create in your domain. Organizational units (OUs) enable you to group objects together so that you can apply a set of policies to the objects. OUs should be designed to group objects by the following criteria:

Object Class Objects can be organized based on their class or type. You can organize all the computers and users into their respective OUs.

Geographic Location Objects can be organized based on their location. It is common practice to use airport codes when the organization is spread across a large distance. However, town names can also be used when the organization branches are in relatively close proximity.

Function Objects can be organized based on their function in the organization, such as servers, workstations, and users. Objects can also be grouped by their job function, such as Sales, Marketing, and HR.

Hybrid Nothing says that you have to stick to one organization of objects. Depending on how you are planning your domain, you could use a combination. It is not uncommon to see airport codes to represent the various offices, with functions under each airport code, as shown in Figure 5.14.

Moving Objects

Moving objects between OUs is a straightforward process that allows administrators to adapt to organizational changes or refine their directory structure. When an object, such as a user or computer, needs to be relocated to a different OU, it can be done through directory management tools like the Active Directory Users and Computers (ADUC) MMC console.

When you move an object from one OU to another, several things happen behind the scenes to ensure the object is properly integrated into its new location:

Policy Application The object will inherit the Group Policy Objects (GPOs) associated with the new OU. This means any policies applied to the new OU, such as security settings, software installations, or login scripts, will now apply to the moved object.


Permissions and Access Control The object will also inherit the permissions and access controls defined for the new OU. This can affect the object’s access to resources and its ability to perform certain actions within the network.

Replication In a multidomain controller environment, the change will be replicated across all domain controllers to ensure consistency. This ensures that the object’s new location is recognized throughout the entire network.

Audit and Logging The move will be logged in the directory service’s audit logs. This helps administrators track changes and maintain a record of object movements for security and troubleshooting purposes.

Script and Automation Adjustments If there are any scripts or automated processes that reference the object’s original location, they may need to be updated to reflect the new OU. This ensures that automated tasks continue to function correctly.

FIGURE 5.14 A hybrid OU structure




Group Policy

Group Policy is a feature of Active Directory that enables you to apply policies to control users and computers. Typically, you do not apply policies to individual users or computers but instead to groups of users or computers. A Group Policy Object (GPO) is a type of object in Active Directory that allows you to apply a set of policies against an organizational

unit. Group Policy Objects are created, linked, and edited in the Group Policy Management Console (GPMC), as shown in Figure 5.15.

FIGURE 5.15 The Group Policy Management Console



You can control thousands of settings for both the user and computer objects, as shown in Figure 5.16. Policies are hard controls that you can force on an object. Policies are refreshed in the background every 90 minutes. So, if a setting that has a policy applied changes, it will be set back during the refresh cycle. Most of the time, however, settings are grayed out when they are being managed by GPO and cannot be changed at all. Preferences allow for files, Registry, environment variable, and Control Panel items to be modified. Preferences set an initial setting and are applied only during first login, so these settings are a preference, not a policy, and the user can change these settings afterward.

Login Scripts

Login scripts are one of the configurable attributes for a user account. As covered in Volume 2 Chapter 8, “Scripting, Remote Access, and Artificial Intelligence,” you can use VBScript or Windows batch scripts as login scripting languages. Login scripts are useful on an Active Directory network for connecting network-mapped drives and printers, among other administrative tasks. Login scripts also provide uniformity across an enterprise by running the same commands for each user configured with the script. The location of the setting is found on the Profile tab of the user account, as shown in Figure 5.17.

FIGURE 5.16 Group Policy Object settings





FIGURE 5.17 Profile settings for a user account



Home Folders

A *home folder* is a private network location in which the user can store their personal files. The home folder is an attribute that can be set for a user account in the Active Directory Users and Computers MMC on the Profile tab, as shown in Figure 5.17. The location can be a local path, if the user will use the same computer, and the files should be stored locally for the user. However, it is most useful when you connect a network drive to a remote file server. This allows for centralized file storage, and you can then perform backups on the data.

Folder Redirection

Normally, when a user logs into the network and a roaming profile exists for the user, the profile is completely downloaded to the computer the user is working on. During logout, all data is written back to the roaming profile location on the network file server. Profiles can become extremely large in size, sometimes even gigabytes, and slow down the login and logout processes.

Folder redirection is a Group Policy setting that allows the redirection of portions of users' profile folders to a network location. When folder redirection is used, the roaming profile is still downloaded. However, the redirected folders are not downloaded; they are simply redirected to the network location. This speeds login and logout times because the entire profile is no longer downloaded (login) and uploaded (logout).

Security Groups

The use of Active Directory domains allows you to adopt a centralized administration model for user accounts. You also have the ability to secure files, printers, and other resources on the domain with these centralized credentials. However, if you secure the resources solely with user accounts, very quickly it will become a daunting task, as more and more people need access to resources. This is a common pitfall of new administrators, because most resources only need to be accessed by a few people initially.

All security should be done using *security groups* for a few reasons. The first reason is a simple one: You want to administer groups of users and not individual users. It is easier to apply permissions to a group of users than individual users. As the resource needs to be shared by more people, you never have to revisit the resource to apply the new permissions if a group is used. All you need to do is add the new users to the group and they will have access.

One additional benefit that accompanies using security groups is the centralized auditing of permissions. If you want to know who has access to a resource, all you have to do is look at the membership of the group associated with the resource. If you need to check if someone specific has permission to the resource, you just have to look at their group membership.

When securing resources, you should always create a new group that is associated with the resource and never reuse a group for multiple resources. For example, if you were securing a main office printer, you should create a group that explains the resource and the level of access, such as creating a group called `perm_print_mainprinter`. Just by looking

at the group you can identify that it is a permissions group, that it allows printing, and that the resource is the main printer.



Exam Objective 2.4, “Summarize types of malware and tools/methods for detection, removal, and prevention,” requires you to know the types of malware. You should know the difference between a Trojan, rootkit, virus, spyware, ransomware, keylogger, boot sector virus, cryptominer, stalkerware, and fileless malware. In addition, you should be familiar with the tools and methods to detect, remove, and prevent the various malware.

Malware

Malware is a broad term describing any software with malicious intent. Although we use the terms *malware* and *virus* interchangeably, distinct differences exist between them. The lines have blurred because the delivery mechanisms of malware and viruses are sometimes indistinguishable.

A virus is a specific type of malware, the purpose of which is to multiply, infect, and do harm. A virus distinguishes itself from other malware because it is self-replicating code that often injects its payload into documents and executables. This is done in an attempt to infect more users and systems. Viruses are so efficient in replicating that their code is often programmed to deactivate after a period of time, or they are programmed to only be active in a certain region of the world.

Malware can be found in a variety of other forms, such as covert cryptomining, web search redirection, adware, spyware, and even ransomware, and these are just a few. Today the largest threat of malware is ransomware because it's lucrative for criminals.

Ransomware

Ransomware is a type of malware that is becoming popular because of anonymous currency, such as Bitcoin. Ransomware is software that is often delivered through an unsuspecting random download, phishing email, malicious attachment, compromised website, or a targeted attack by a threat actor. Once on a computer or in the network, it takes control of systems and demands that a third party be paid. The “control” can be accomplished by encrypting the hard drive, by changing user password information, or via any of a number of other creative ways. Users are usually assured that by paying the extortion amount (the ransom), they will be given the code needed to revert their systems back to normal operations. CryptoLocker was a popular ransomware that made headlines across the world (see Figure 5.18), but there have been many others since its discovery. You can protect yourself from ransomware by having antivirus/anti-malware software with up-to-date definitions and by keeping current on patches.

FIGURE 5.18 CryptoLocker

Trojans

Trojan horses are programs that enter a system or network under the guise of another program. A Trojan horse may be included as an attachment or as part of an installation program. The Trojan horse can create a backdoor or replace a valid program during installation. It then accomplishes its mission under the guise of another program. Trojan horses can be used to compromise the security of your system, and they can exist on a system for years before they're detected.

The best preventive measure for Trojan horses is to not allow them entry into your system. Immediately before and after you install a new software program or operating system, back it up! If you suspect a Trojan horse, you can reinstall the original program(s), which should delete the Trojan horse. A port scan may also reveal a Trojan horse on your system. If an application opens a TCP or UDP port that isn't supported in your network, you can track it down and determine which port is being used.

Keyloggers

A *keylogger* is normally a piece of software that records an unsuspecting victim's keystrokes. Keyloggers can stay loaded in memory and wait until you log into a website or other authentication system. They will then capture and relay the information to an awaiting host on the Internet.

Keyloggers don't always have to be in the form of software. Some keyloggers are hardware dongles that sit between the keyboard and computer. These must be retrieved and the data must be downloaded manually, so they are not very common.

Rootkits

Rootkits are software programs that have the ability to hide certain things from the operating system. They do so by obtaining (and retaining) administrative-level access. With a rootkit, there may be a number of processes running on a system that don't show up in Task Manager, or connections that don't appear in a Netstat display of active network connections that may be established or available. The rootkit masks the presence of these items by manipulating function calls to the operating system and filtering out information that would normally appear.

Unfortunately, many rootkits are written to get around antivirus and antispyware programs that aren't kept up-to-date. The best defense you have is to monitor what your system is doing and catch the rootkit in the process of installation.

Spyware

Spyware differs from other malware in that it works—often actively—on behalf of a third party. Rather than self-replicating, like viruses and worms, spyware is spread to machines by users who inadvertently ask for it. The users often don't know they have asked for it but have done so by downloading other programs, visiting infected sites, and so on.

The spyware program monitors the user's activity and responds by offering unsolicited pop-up advertisements (sometimes known as *adware*), gathers information about the user to pass on to marketers, or intercepts personal data, such as credit card numbers.

Cryptominers

With the rise of bitcoin, so came the rise of cryptominers. A cryptominer is typically a purpose-built device that grinds out cryptographic computations. When the computation is balanced, a cryptocurrency is created and equates to real money, such as Bitcoin, Ethereum, and Dogecoin, just to name a few. A cryptominer does not always have to be a dedicated purpose-built device—it can also be a distributed group of computers called a cryptopool.

Malware in the form of cryptominers became very popular, because it is a very lucrative way for threat agents to make money. The problem is that the threat agents use your computer to grind out the computations. The most common way a threat agent will run a cryptominer remotely is with JavaScript embedded on a malicious web page. Threat agents have also been known to create viruses in which the payload (cryptominer) uses your video card to grind out the computations. However, the JavaScript variant is more common to find in the wild.

Fileless

Fileless malware is a sophisticated type of cyberthreat that operates without leaving a trace on the victim's hard drive, making it exceptionally difficult to detect and remove. Unlike

traditional malware, which relies on executable files to infect a system, fileless malware exploits existing software, applications, and protocols to execute malicious activities directly in the system's memory. This stealthy approach allows it to bypass conventional antivirus solutions and persist in the system undetected. Cybercriminals often use fileless malware to gain unauthorized access, steal sensitive information, or disrupt operations, posing a significant challenge for cybersecurity professionals. As technology evolves, so do the tactics of these malicious actors, necessitating advanced and adaptive security measures to combat such elusive threats.

Stalkerware

Stalkerware is a particularly insidious form of spyware designed to monitor and track an individual's activities without their knowledge or consent. Often disguised as legitimate applications, stalkerware can be secretly installed on a victim's device, granting the perpetrator access to personal data such as messages, call logs, location, and even real-time audio and video feeds. This invasive software is frequently used in abusive relationships to exert control and surveillance over the victim, leading to severe privacy violations and emotional distress. The covert nature of stalkerware makes it challenging to detect and remove, highlighting the need for robust cybersecurity measures and increased awareness to protect individuals from such intrusive threats. As technology continues to advance, it is crucial to stay vigilant and informed about the potential risks posed by stalkerware and other forms of digital surveillance.

Viruses

Viruses can be classified as polymorphic, stealth, retrovirus, multipartite, armored, companion, phage, and macro viruses. Each type of virus has a different attack strategy and different consequences.

The following sections introduce the symptoms of a virus infection, explain how a virus works, and describe the types of viruses you can expect to encounter and how they generally behave. We'll also discuss how a virus is transmitted through a network and look at a few hoaxes.



Estimates for losses due to viruses are in the billions of dollars. These losses include financial loss as well as lost productivity.

Symptoms of a Virus/Malware Infection

Many viruses will announce that you're infected as soon as they gain access to your system. They may take control of your system and flash annoying messages on your screen or destroy your hard disk. When this occurs, you'll know that you're a victim. Other viruses will cause your system to slow down, cause files to disappear from your computer, or take over your disk space.



Because viruses are the most common type of malware, the term *virus* is used in this section.

You should look for some of the following symptoms when determining if a virus infection has occurred:

- The programs on your system start to load more slowly. This happens because the virus is spreading to other files in your system or is taking over system resources.
- Unusual files appear on your hard drive, or files start to disappear from your system. Many viruses delete key files in your system to render it inoperable.
- Program sizes change from the installed versions. This occurs because the virus is attaching itself to these programs on your disk.
- Your browser, word processing application, or other software begins to exhibit unusual operating characteristics. Screens or menus may change.
- The system mysteriously shuts itself down or starts itself up and does a great deal of unanticipated disk activity.
- You mysteriously lose access to a disk drive or other system resources. The virus has changed the settings on a device to make it unusable.
- Your system suddenly doesn't reboot or gives unexpected error messages during startup.

This list is by no means comprehensive. What is an absolute, however, is the fact that you should immediately quarantine the infected system. It is imperative that you do all you can to contain the virus and keep it from spreading to other systems within your network, or beyond.


How Viruses Work

A virus, in most cases, tries to accomplish one of two things: render your system inoperable or spread to other systems. Many viruses will spread to other systems given the chance and then render your system unusable. This is common with many of the newer viruses.

If your system is infected, the virus may try to attach itself to every file in your system and spread each time you send a file or document to other users. Some viruses spread by infecting files that are either transmitted through a network or by removable media, such as backup tapes, USB thumb drives, among others. When you give removable media to another user or put it into another system, you then infect that system with the virus.

Many viruses today are spread using email. The infected system attaches a file to any email that you send to another user. The recipient opens this file, thinking it's something that you legitimately sent them. When they open the file, the virus infects the target system. The virus might then attach itself to all the emails that the newly infected system sends, which in turn infects the computers of the recipients of the emails. Figure 5.19 shows how a virus can spread from a single user to literally thousands of users in a very short period of time using email.

FIGURE 5.19 A virus spreading from an infected system using email



Types of Viruses

Viruses take many different forms. The following list briefly introduces these forms and explains how they work.



NOTE

The best defense against a virus attack is up-to-date antivirus/antimalware software installed and running. The software should be on all workstations as well as the server.

These are the most common types of viruses, but this isn't a comprehensive list:

Armored Virus An armored virus is designed to make itself difficult to detect or analyze. Armored viruses cover themselves with protective code that stops debuggers or disassemblers from examining critical elements of the virus. The virus may be written in such a way that some aspects of the programming act as a decoy to distract analysis while the actual code hides in other areas in the program.

From the perspective of the creator, the more time that it takes to deconstruct the virus, the longer it can live. The longer it can live, the more time it has to replicate and spread to as many machines as possible. The key to stopping most viruses is to identify them quickly and educate administrators about them—the very things that the armor makes difficult to accomplish.


Boot Sector Virus This virus infects the Master Boot Record (MBR) of a hard disk. This type of virus loads when the computer boots and can re-infect an operating system. Secure Boot is used to verify the entire boot process with digital signatures and identify any part of the boot process that has been modified.

Companion Virus A companion virus attaches itself to legitimate programs and then creates a program with a different filename extension. This file may reside in your system's temporary directory. When a user types the name of the legitimate program, the companion virus executes instead of the real program. This effectively hides the virus from the user. Many of the viruses that are used to attack Windows systems make changes to program pointers in the Windows Registry so that they point to the infected program. The infected program may perform its dirty deed and then start the real program.

Macro Virus A macro virus exploits the enhancements made to many application programs. Programmers can expand the capability of applications such as Microsoft Word and Excel. Word, for example, supports a mini-BASIC programming language that allows files to be manipulated automatically. These programs in the document are called *macros*. For example, a macro can tell your word processor to spell-check your document automatically when it opens. Macro viruses can infect all the documents on your system and spread to other systems via email or other methods. Macro viruses are one of the fastest-growing forms of exploitation today.

Multipartite Virus A multipartite virus attacks your system in multiple ways. It may attempt to infect your boot sector, infect all your executable files, and destroy your application files. The hope here is that you won't be able to correct all the problems and will allow the infestation to continue. The multipartite virus depicted in Figure 5.20 attacks a system's boot sector, infects application files, and attacks Word documents.

FIGURE 5.20 A multipartite virus commencing an attack on a system



Phage Virus A phage virus alters programs and databases, and the only way to remove this virus is to reinstall the programs that are infected. If you miss even a single instance

of this virus on the victim system, the process will start again and infect the system once more.

Polymorphic Virus Polymorphic viruses change form to avoid detection. These types of viruses attack your system, display a message on your computer, and delete files on your system. The virus will attempt to hide from your antivirus software. Frequently, the virus will encrypt parts of itself to avoid detection. When the virus does this, it's referred to as mutation. The mutation process makes it hard for antivirus software to detect common characteristics of the virus.



A *signature* is an algorithm or other element of a virus that uniquely identifies it. Because some viruses have the ability to alter their signature, it is crucial that you keep signature files current, whether you choose to download them manually or configure the antivirus engine to do so automatically.

Retrovirus A *retrovirus attack* bypasses the antivirus software installed on a computer. You can consider a retrovirus to be an anti-antivirus. Retroviruses can directly attack your antivirus software and potentially destroy the virus definition database file. When this information is destroyed without your knowledge, you are left with a false sense of security. The virus may also directly attack an antivirus program to create bypasses for itself.

Stealth Virus A *stealth virus* attempts to avoid detection by masking itself from applications. It may attach itself to the boot sector of the hard drive. When a system utility or program runs, the stealth virus redirects commands around itself to avoid detection. An infected file may report a file size different from what is actually present. Stealth viruses may also move themselves from file A to file B during a virus scan for the same reason.

Virus Transmission in a Network

Upon infection, some viruses destroy the target system immediately. The saving grace is that the infection can be detected and corrected. Some viruses won't destroy or otherwise tamper with a system; instead, they use the victim system as a carrier. The victim system then infects servers, fileshares, and other resources with the virus. The carrier then infects the target system again. Until the carrier is identified and cleaned, the virus continues to harass systems in this network and spread.

Present Virus Activity

New viruses and threats are released on a regular basis to join the cadre of those already in existence. From an exam perspective, you need only be familiar with the world as it existed at the time the questions were written. From an administrative standpoint, however, you need to know what is happening today.

To find information on alerts and threats, visit the Cybersecurity and Infrastructure Security Agency (CISA) Cybersecurity Alerts & Advisories web page at www.cisa.gov/news-events/cybersecurity-advisories. You'll find a detailed description of the most current threats as well as links to pages on older threats and advisories of vulnerable systems.

Botnets

A *botnet* is a group of zombies, which sounds like a ridiculous beginning to a horror movie. When malware infects a computer, its purpose is often to lie dormant and await a command from a command-and-control server. When this happens, the computer is considered a zombie. When enough infected computers (zombies) check in, the *threat agent* will send a command to the command-and-control server, and the botnet of zombies will work on the task. Often the task is to launch a malicious DDoS attack or to send spam. DDoS will be covered later in this chapter, in the “Common Security Threats” section.

Worms

A *worm* is different from a virus in that it can reproduce itself, it's self-contained, and it doesn't need a host application to be transported. Many of the so-called viruses that make the news are actually worms. However, it's possible for a worm to contain or deliver a virus to a target system.

By their nature and origin, worms are supposed to propagate, and they use whatever services they're capable of using to do that. Early worms filled up memory and bred inside the RAM of the target computer. Worms can use TCP/IP, email, Internet services, or any number of possibilities to reach their target.

Adware

Adware is a type of software that automatically delivers advertisements to a user's device, often without their consent. These ads can appear in various forms, such as pop-ups, banners, or in-browser ads, and are typically designed to generate revenue for the software's creator. While adware itself is not inherently malicious, it can significantly degrade the user experience by slowing down the device, consuming bandwidth, and creating a cluttered and intrusive environment. Additionally, some adware may track user behavior and collect personal information, raising privacy concerns. To mitigate the impact of adware, users should be cautious when downloading free software, as adware is often bundled with these applications, and should use reputable security tools to detect and remove unwanted programs.

Potentially Unwanted Programs (PUPs)

Potentially unwanted programs (PUPs) are software applications that a user may not have intended to install on their device. These programs often come bundled with legitimate

software and can include adware, toolbars, or other unwanted utilities. While PUPs are not always harmful, they can negatively affect system performance, compromise privacy, and create security vulnerabilities. PUPs can be particularly deceptive, as they often masquerade as useful tools or enhancements, tricking users into installing them. To avoid PUPs, users should carefully read installation prompts, opt for custom installation options to deselect unwanted components, and regularly scan their devices with trusted antimalware software to identify and remove any PUPs that may have been inadvertently installed.

Mitigating Software Threats

Now that you understand some of the common software threats, let's look at how you can protect yourself from them. This section discusses practical tools and methods that you can use to safeguard yourself from common software-based threats. We will also discuss tactics you can use to mitigate risk as well as ways to recover from security mishaps.


Antivirus

Most malware can be simply prevented with the use of antivirus software. Back when Windows XP came out, the running joke was that you would get a virus before you could get a chance to install antivirus software. To some extent this was true, if you had to get online to retrieve the software.

Microsoft introduced Microsoft Security Essentials as a download for Windows XP, and the Windows Vista operating system started to ship with it installed. Today, Windows comes preinstalled with Windows Virus & Threat Protection, so if you don't purchase antivirus software you are still protected. As a result of these tactics, Microsoft has made the Windows operating system safer than it used to be.

Although Microsoft's antivirus program will work fine for most computing needs, there are some advantages to purchasing antivirus products from third-party vendors. To understand some of the differences, you need to be familiar with the components of antivirus software. Antivirus software comprises two main components: the antivirus engine and the definitions database, as shown in Figure 5.21.

FIGURE 5.21 Antivirus software components



Antivirus Engine The antivirus engine is responsible for the real-time scanning of operating system files and the notifications to the user. Various antivirus engines will scan operating system files differently. For example, an antivirus engine might recognize when a game is being played, and it will temporarily stop scanning files so that the game loads faster. Notifications can also automatically submit files to the cloud for more extensive scanning.


Definitions Database The definitions database is why you would select one antivirus product over another. A definitions database's effectiveness is based on two attributes: the frequency of updates and the comprehensiveness of the database signatures. Antivirus definitions, also known as antivirus signatures, are discovered daily and added to the antivirus database. Some antivirus products will download deltas to the initial database every hour, and some will download them once a day.

Antimalware

Malware is a broad term and covers many different software threats that we learned about in this chapter. Antimalware and antivirus are extremely similar in their functionality, and sometimes vendors have a hard time differentiating their products. This is because many antivirus products now check more than just files.

An antimalware software package will not only check the filesystem for threats, like root-kits and trojans, but will also watch incoming email for phishing scams and malicious web-sites, as shown in Figure 5.22. When these threats are detected, the user gets a notification, and the threat is usually mitigated or avoided completely.

FIGURE 5.22 Antimalware components



In Exercise 5.1 you will test your antimalware protection with a harmless file called an Eicar file.

EXERCISE 5.1**Testing Your Antimalware**

1. Navigate to the Eicar antimalware test file site at www.eicar.org/download-anti-malware-testfile.
2. Scroll down to the download section.
3. Download a few of the Eicar test files and notice how your antivirus detects the malware.
4. Examine the alerts as your antimalware software reports the malware.

The Eicar website contains a totally benign piece of malware that triggers your antimalware engine. Any search for Eicar will produce similar results and the contents are benign.

Detection and Response

The increasing sophistication and frequency of cyberthreats necessitate advanced detection and response mechanisms beyond traditional antivirus (AV) or antimalware solutions. While AV and antimalware software are essential for identifying and blocking known threats, they often fall short in detecting and responding to more complex and evolving attacks. This is where endpoint detection and response (EDR), managed detection and response (MDR), and extended detection and response (XDR) come into play. These advanced solutions offer distinct capabilities and advantages, such as continuous monitoring, threat hunting, and comprehensive visibility across multiple security layers, making them crucial for addressing specific security challenges and requirements in today's dynamic threat landscape. The following explains how each of these detection and response packages differentiate themselves from each other:

Endpoint Detection and Response (EDR) Endpoint detection and response (EDR) is a cybersecurity solution designed to monitor and protect endpoints, such as computers, mobile devices, and servers, from malicious activities. EDR solutions continuously collect and analyze data from endpoints to detect suspicious behavior, enabling rapid identification and response to potential threats. By providing detailed visibility into endpoint activities, EDR assists security teams in quickly investigating and mitigating incidents, thereby reducing the risk of data breaches and minimizing the impact of cyberattacks.

Managed Detection and Response (MDR) Managed detection and response (MDR) is a comprehensive cybersecurity service that combines advanced technology with human expertise to detect, analyze, and respond to threats. MDR providers offer 24/7 monitoring and threat hunting, leveraging sophisticated tools and techniques to identify and neutralize threats before they can cause significant harm. By outsourcing these critical functions to a team of skilled security professionals, organizations can enhance their

security posture without the need for extensive in-house resources. MDR services are particularly beneficial for small to medium-sized businesses that may lack the expertise or budget to maintain a full-fledged security operations center (SOC).

Extended Detection and Response (XDR) Extended detection and response (XDR) is an integrated cybersecurity solution that extends beyond traditional EDR by correlating data across multiple security layers, including endpoints, networks, servers, and cloud environments. XDR provides a holistic view of an organization's security landscape, enabling more effective threat detection and response. By breaking down silos and integrating various security tools and data sources, XDR enhances the ability to detect complex, multivector attacks and streamline incident response processes. This comprehensive approach helps organizations improve their overall security posture and reduce the time and effort required to manage and respond to threats.

Each detection and response platform plays a vital role in protecting an organization from compromise. EDR focuses on endpoint protection, providing detailed visibility and rapid response capabilities. MDR offers a managed service approach, combining technology and expertise to deliver continuous threat monitoring and response. XDR takes a broader view, integrating data from multiple sources to provide a comprehensive security solution. By understanding and leveraging these products, organizations can better protect themselves against cyberthreats, and ensure the safety and integrity of their data.

Email Security Gateway Systems

Email security gateway systems are a critical component of an organization's cybersecurity defenses. They are designed to protect against a number of email threats such as phishing, malware, ransomware, and of course spam. These systems function as a protective layer, inspecting incoming and outgoing email traffic to filter out malicious content and prevent it from reaching end users. Email security gateways can be deployed in two primary forms: on-premises and cloud-based systems.

On-Premises Email Security Gateways On-premises email security gateways are installed and managed within the organization's own data centers. This deployment model provides organizations with direct control over their security policies, configurations, and data. On-premises solutions are often favored by organizations with stringent data privacy requirements, such as those in the healthcare or financial sectors, or those that need to comply with specific regulatory standards like GDPR or HIPAA. However, on-premises solutions require significant investment in hardware, software, and ongoing maintenance, as well as dedicated IT staff to manage and update the system.

Cloud-Based Email Security Gateways Cloud-based email security gateways, on the other hand, are hosted by third-party providers and accessed over the Internet. This deployment model offers several advantages, including scalability, flexibility, and reduced maintenance overhead. Cloud-based solutions can easily scale to accommodate the growing needs of an organization, making them ideal for businesses of all sizes.


They also benefit from the provider's advanced threat intelligence, which is gathered from a global network of users and continuously updated to protect against the latest threats. Additionally, cloud-based gateways often come with built-in redundancy and disaster recovery capabilities, ensuring continuous protection even in the event of a system failure. Organizations with limited IT resources or those looking to quickly adapt to evolving threat landscapes may find cloud-based solutions particularly beneficial. However, reliance on a third-party provider means that organizations must trust the provider's security measures and data handling practices.

Both on-premises and cloud-based email security gateways play a crucial role in safeguarding an organization from email compromise. They filter out malicious content and prevent data breaches, thus protecting against financial loss, and ensure compliance with regulatory requirements.

Recovery Console


A recovery console can perform a number of useful functions for recovery from a security threat. The Windows Recovery Environment (WinRE) is a recovery console that can perform a number of useful functions, as we'll cover in this section. The most useful function is the Reset This PC option, which allows you to refresh the operating system while keeping your data files or remove everything and start from scratch, as shown in Figure 5.23. The latter of the two options assumes you have backups of your data files.

FIGURE 5.23 Windows Recovery Environment



The Windows Recovery Environment also allows you to perform a system restore, whereby you can restore the operating system back to a specific point in time. If a system recovery image exists, you can also recover with the System Image Recovery option. This option will reset the operating system back to the point in the recovery image, which is usually just like the day you turned it on. Figure 5.24 shows the Advanced Options menu.

FIGURE 5.24 Windows Advanced Options



To access Windows Recovery Advanced Options, go to Start > Settings > System > Recovery, then under Advanced Startup, select Restart Now; this will boot you into the Windows Recovery Environment where you can choose Troubleshoot and then Advanced Options to access further recovery tools.

End-User Education

By far the best prevention of security threats is the education of your end users regarding common threats. For example, the most effective method of preventing viruses, malware, spyware, and harm to data is to teach your users not to open suspicious files and to open only those files that they're reasonably sure are virus/malware free. End users should also

be educated on how to identify Trojans and phishing emails scams. The end-user education should also identify guidelines for physical destruction of data, in particular any paperwork that has sensitive information on it. End users should also be educated on the various social engineering threats and how to identify them. An end user who has foresight and who exercises vigilance is more powerful than any antivirus or antimalware product on the market.

End-user education in an organization is normally part of the employee onboarding process for new hires. However, it should not stop there, because threats change every day. Many organizations revisit the training for their employees once a year in ideal circumstances. This training can be performed in a formal classroom setting or through an online service. Some online services offer educational videos that have interactive questions to verify that the employee has learned the objectives of the video.

Because phishing is such a widespread problem for organizations, special antiphishing training is often mandatory for employees a few times a year. Often organizations will phish their employees with specially crafted emails in an attempt to see how well their training is working. When an employee spots the phishing attempt, they can earn rewards, like a gift card. However, if they get phished, then they must retake the antiphishing training or they may be targeted in the future. A popular month for these tactics is October, because it is the cybersecurity awareness month.


Software Firewalls

“Software firewalls” is a misnomer for this section, since all firewalls are software-based in some way. Sure, you might purchase a piece of equipment that is classified as a hardware firewall, but there is software running on the firewall to protect your network. However, when we discuss firewalls in respect to operating systems, we call them software firewalls, because they are part of the operating system and thus considered software.

Let’s look at where we came from and where we are now. A major gamechanger in the history of Microsoft was the release of Windows XP Service Pack 2, which switched on the built-in firewall by default. It was long overdue for the operating system at the time. It was also not well received, because administrators had to learn firewall rules when they installed a new software package that required incoming network traffic. As a result, the firewall was the first thing that got shut off when there was a problem with connectivity.


Throughout the years of new Windows versions, the firewall has received new features and became a polished product. Windows 10 renamed the product Windows Defender Firewall and Windows Defender Firewall with Advanced Security, as shown in Figure 5.25. Both of these configure the same firewall service, with the latter of the two allowing for much more granular control. The inner workings are almost identical to those of the original firewall that shipped with Windows XP SP2.

The Windows firewall has achieved its original purpose of protecting the operating system from malicious worms and malicious inbound network connections. By default, the outbound network traffic is allowed and inbound network traffic is blocked, unless a rule exists, as shown in Figure 5.26.

FIGURE 5.25 Windows Defender Firewall with Advanced Security on Windows 11

You will also notice that there are three different profiles: Domain, Private, and Public. When the network service starts up, it contacts the default gateway (router) and configures itself to a profile. This allows the operating system to be location-aware and protect itself differently based on your location. In prior versions of Windows, if the router has never been seen before, then you'll get prompted with a dialog box asking you to choose if you want to allow your PC to be discoverable. If you answer Yes, then the firewall profile will be configured as Private, and any rules associated with the Private Profile will be active. If you answer No, then any rules associated with the Public Profile will be active. In Windows 11, the location profile will automatically default to the Public Profile. If you want to change the profile, you can navigate to Start menu ► Settings ► Network & Internet ► Properties of the current connection, then Public Network or Private Network. The Domain Profile is automatically selected if the network is the corporate network and the operating system is joined to the domain.

Firewalls are also built into other operating systems, including Linux and macOS. Depending on the distribution of Linux or macOS, the firewall included will vary, as well as the way you would configure it. However, most distributions of Linux, such as Ubuntu and Debian, come with the iptables firewall installed. Fedora Linux comes with firewalld, which also supports location-based firewall rules.

FIGURE 5.26 Windows Defender Firewall with Advanced Security defaults

Regardless of which type of firewall an operating system comes preinstalled with, a third-party firewall can be installed. These firewalls can offer intrusion-detection capabilities that alert you when someone is attacking. In almost all cases, the firewall that comes with the operating system is more than adequate, but you have to keep it on at all times to prevent unwanted connections.

Reinstalling the OS

When you are compromised by a virus or other type of malware, the only way to be sure you have removed it completely is to reinstall the operating system. This may seem like an extreme measure, but virus researchers do not always know what the threat agent embeds in the operating system. The threat agent's mission is to gain access to your operating system and to keep a persistent connection. If that means opening a few other backdoors, then that is what they will embed in their malware.

Fortunately, the Windows operating system makes it easy to reinstall the operating system. The Windows operating system allows you to reset the PC with the use of the recovery console or from the Settings app. Many devices also have a recovery tool embedded in the browser to factory-reset/reimage the device.



Exam Objective 2.5, “Compare and contrast common social engineering attacks, threats, and vulnerabilities,” requires you to know the types of social engineering. You should be familiar with all of the different social engineering types, threats, vulnerabilities, and the tactics a threat actor will take to compromise your organization with social engineering.

Social Engineering Attacks, Threats, and Vulnerabilities

Social engineering is a process in which an attacker attempts to acquire information about your network and system by social means, such as talking to people in the organization. A social engineering attack may occur over the phone, by email, or in person. The intent is to acquire access information, such as user IDs and passwords. When the attempt is made through email or instant messaging, it is known as *phishing* (discussed later), and it's often made to look as if a message is coming from sites where users are likely to have accounts. (Banks, bills, and credit cards are popular.)

These are relatively low-tech attacks and are more akin to con jobs. Take the following example: Your help desk gets a call at 4:00 a.m. from someone purporting to be the vice president of your company. They tell the help desk personnel that they are out of town to attend a meeting, their computer just failed, and they are sitting in a FedEx office trying to get a file from their desktop computer back at the office. They can't seem to remember their password and user ID. They tell the help desk representative that they need access to the information right away or the company could lose millions of dollars. Your help desk rep knows how important this meeting is and gives the user ID and password over the phone. At this point, the attacker has just successfully socially engineered an ID and password that can be used for an attack by impersonating a high-profile person.

Another common approach is initiated by a phone call or email from someone who pretends to be your software vendor, telling you that they have a critical fix that must be installed on your computer system. It may state that if this patch isn't installed right away, your system will crash and you'll lose all your data. For some reason, you've changed your maintenance account password, and they can't log in. Your system operator gives the password to the person. You've been hit again.

In Exercise 5.2, you'll test your users to determine the likelihood of a social engineering attack. The steps are suggestions for tests; you may need to modify them slightly to be appropriate at your workplace. Before proceeding, make certain that your manager knows that you're conducting such a test and approves of it.

EXERCISE 5.2

Testing Social Engineering

1. Call the receptionist from an outside line when the sales manager is at lunch. Tell the receptionist that you're a new salesperson, that you didn't write down the username and password the sales manager gave you last week, and that you need to get a file from the email system for a presentation tomorrow. Does the receptionist direct you to the appropriate person or attempt to help you retrieve the file?

2. Call the human resources department from an outside line. Don't give your real name but instead say that you're a vendor who has been working with this company for years. You'd like a copy of the employee phone list to be emailed to you, if possible. Do they agree to send you the list, which would contain information that could be used to try to guess usernames and passwords?
3. Pick a user at random. Call them and identify yourself as someone who works with the company. Tell them that you're supposed to have some new software ready for them by next week and that you need to know their password to finish configuring it. Do they do the right thing?

The best defense against any social engineering attack is education. Make certain that the employees of your company know how to react to the requests presented here. Social engineering works on the premise that people try to help when they are vested in your efforts, such as a co-worker or if you are trying to help them.

Phishing

Phishing is a form of social engineering in which you ask someone for a piece of information that you are missing by making it look as if it is a legitimate request. An email might look as if it is from a bank and contain some basic information, such as the user's name. These types of messages often state that there is a problem with the person's account or access privileges. The person will be told to click a link to correct the problem. After they click the link, which goes to a site other than the bank's, they are asked for their username, password, account information, and so on. The person instigating the phishing attack can then use this information to access the legitimate account.



One of the best countermeasures to phishing is to mouse over the Click Here link and read the URL. Almost every time, the URL is an adaptation of the legitimate URL, as opposed to a link to the real thing.

The only preventive measure in dealing with social engineering attacks is to educate your users and staff never to give out passwords and user IDs over the phone or via email or to anyone who isn't positively verified as being who they say they are.

Vishing

When phishing is combined with *Voice Over IP (VoIP)*, it becomes known as *vishing*, which is just an elevated form of social engineering. While crank calls have existed since the invention of the telephone, the rise in VoIP now makes it possible for someone to call you from almost anywhere in the world, without the worry of tracing, caller ID, and other features of landlines, and pretend to be someone they are not in order to get data from you.

Smishing

Smishing is the same thing as vishing and phishing, except they text it to your phone instead with a Short Message Service (SMS) code. The message typically has an emotional appeal, urgency, and call to action. I currently get about 10 texts a day that range from saying my packages can't be delivered and I must call, my online services are being shut off for nonpayment, or the government has decided to give me money! In fact, I looked at my phone while writing this and found this message:

Netflix: We couldn't process your monthly payment, so your account is temporarily on hold. Please update your billing information here: <https://Netflix-Canada-update.com>.

Typically the threat actors will want iTunes gift cards and these can then be turned into money. That or the threat actors really dig music, but I think the motivation is money.

Quishing

QR code phishing, also known as *quishing*, is a sophisticated social engineering tactic where cybercriminals use QR codes to lure victims into scanning them with their mobile devices. Once scanned, these QR codes redirect users to malicious websites designed to steal sensitive information, such as login credentials or financial data. Attackers often embed these QR codes in seemingly legitimate emails or PDF documents, leveraging brand impersonation and urgency to trick users into scanning them. This method is particularly effective because traditional email filters struggle to detect QR codes, and the use of multiple devices, such as scanning with a personal phone can bypass corporate security measures. As a result, quishing poses a unique challenge for both individuals and organizations in maintaining cybersecurity. Don't scan random QR codes on gas station pumps when you are pumping gas; they might be malicious in nature.

Spear Phishing

With spear phishing, the attacker uses information that the target would be less likely to question because it appears to be coming from a trusted source. Suppose, for example, that you receive a message that appears to be from your spouse that says to click here to see that video of your children from last Christmas. Because it appears far more likely to be a legitimate message, it cuts through your standard defenses like a spear, and the likelihood that you would click this link is higher. Generating the attack requires much more work on the part of the attacker, and it often involves using information from contact lists, friend lists from social media sites, and so on.



Trust is a key issue with security. There are trusted software sources that you know and work with all the time (such as Microsoft), and there are untrusted sources—and you should differentiate between them. Likewise, there are trusted sites and untrusted sites, and you should similarly differentiate between them. Let common sense be your guide.

Whaling

Whaling is nothing more than phishing, or spear phishing, for so-called “big” users—thus, the reference to the ocean’s largest creatures. Instead of sending out a To Whom It May Concern message to thousands of users, the whaler identifies one person from whom they can gain all the data that they want—usually a manager or business owner—and targets the phishing campaign at them.

Shoulder Surfing

Another form of social engineering is known as *shoulder surfing*. It involves nothing more than watching someone when they enter their sensitive data. They can see you entering a password, typing in a credit card number, or entering any other pertinent information. A privacy filter can be used to block people from looking at your screen from an angle. However, privacy filters do not protect you as you are entering a password, since a shoulder surfer will watch your keystrokes. The best defense against this type of attack is to survey your environment before entering personal data. It is also proper etiquette to look away when someone is entering their password.

Tailgating

Tailgating is another form of social engineering, and it works because we want to be helpful. Tailgating is the act of entering a building that requires a swipe card or other authentication factor by using the person in front of you.

You may be walking toward an entry that requires some authentication, when someone walking the same way introduces themselves as new to the company and shares some stories about their first day. By the time you get to the door, you may hold it open for them and wish them luck. It can even happen without you knowing it, if the door barely closes and they grasp it. Several different tactics, such as access control vestibules and guards, can be used to mitigate this threat. The best prevention is education of your staff to make sure that it does not happen.

Impersonation

Impersonation is prevalent in many different social engineering attacks. The threat agent portrays (impersonates) another employee for many of the attacks to work. Most employees want to help another fellow employee. The threat agent might be impersonating Bill Jones from accounting asking IT to reset his password. Or the threat agent might impersonate the IT department calling Bill Jones and instructing him to type some well-crafted commands into his computer. Many phishing emails also impersonate your bank, an online store, or some other reputable source in an attempt to steal your credentials. The best method for combatting impersonation is end-user training. Training should help users identify suspicious email or phone calls.

Dumpster Diving

Dumpster diving is the act of a person rifling through the trash with the expectation to find information. A strong policy to prevent dumpster diving is the physical destruction of any sensitive data. Destruction can be performed with the use of a mechanical shredder on-site or a service that destroys materials off-site on behalf of the organization. If an off-site shredding service is performed, always request a signed certificate of destruction to prove that sensitive material was destroyed.

Common Security Threats



A *threat* is a potential danger to the network or the assets of the organization. The potential danger to a network or organization is the attack that a threat agent can carry out. All attacks upon an organization are either technology based or physically based. A technology-based attack is one in which the network and operating systems are used against the organization in a negative way. Physically based attacks use human interaction or physical access, which we previously covered as social engineering attacks. We will now cover several different types of technology-based attacks that are commonly used against networks and organizations.

Denial-of-Service Attacks


A *denial-of-service (DoS)* is an attack launched to disrupt the service or services a company receives or provides via the Internet. A DoS attack is executed with an extremely large number of false requests; because of the attack, the servers will not be able to fulfill valid requests for clients and employees. There are several different types of DoS attacks:

Reflective A reflective DoS is not a direct attack; it requires a third party that will inadvertently execute the DoS. The attacker will send a request to a third-party server and forge the source address of the packet with the victim's IP address. When the third party responds, it responds back to the victim. There are two victims in this type of DoS attack: the first is the victim the attack is aimed at, and the second is the third-party server used to carry out the attack. Figure 5.27 shows an ICMP-based smurf attack, which is a reflective attack.

Amplified An amplified DoS is a variant of a reflective DoS attack. It is carried out by making a small request to the third-party server that yields a larger response to the victim. The most common third-party servers used to carry out the attack are DNS and NTP. For example, an attacker will request a DNS query for a single hostname that contains 20 aliases while forging the source IP address. The victim is then barraged with the 20 answers from the query, as shown in Figure 5.28.


FIGURE 5.27 An ICMP-based smurf attack**FIGURE 5.28** An amplified attack

Distributed A distributed denial-of-service (DDoS) has become the most common type of DoS, because the source of the DoS is varied. It is common for botnets to launch DDoS attacks on organizations. When a single host is used to create a DoS, it can simply be blocked. However, when traffic is coming from millions of different hosts, it is impossible to isolate the DoS and firewall the source. Figure 5.29 shows a DDoS attack.

FIGURE 5.29 A DDoS attack

Evil Twin

An *evil twin* attack is a wireless phishing attack in which the attacker sets up a wireless access point to mimic the organization's wireless access points. When a user connects to the evil twin, it allows the attacker to listen in on the user's traffic. Evil twin access points often report a stronger signal to entice the user to connect to the specific access point, as shown in Figure 5.30. The attacker will then create a connection back to the wireless network and passively sniff network traffic as it routes the traffic to the original destination. The best way to mitigate against evil twin attacks is to perform wireless site surveys on a regular basis to ensure that only valid access points are being used.

FIGURE 5.30 Evil twin attack

Zero-Day Attacks

When a hole (vulnerability) is found in a web browser or other software, and attackers begin exploiting it the very day it is discovered by the developer (bypassing the one-to-two-day response time that many software providers need to put out a patch once the hole has been found), it is known as a *zero-day* attack (or exploit). It is very difficult to respond to a zero-day exploit. If attackers learn of the weakness the same day as the developer, then they have the ability to exploit it until a patch is released. Often, the only thing that you as a security administrator can do, between the discovery of the exploit and the release of the patch, is to turn off the service. You can do this by isolating or disconnecting the system(s) from the network until a patch is released. Although this can be a costly undertaking in terms of productivity, it is the only way to keep the network safe.



Several years ago, Stuxnet was found to be using a total of four zero-day vulnerabilities to spread from host to host. You can learn more about Stuxnet here:

www.kaspersky.com/resource-center/definitions/what-is-stuxnet


Spoofing Attacks

A *spoofing attack* is an attempt by someone or something to masquerade as someone else. This type of attack is usually considered an access attack. A common spoofing attack that was popular for many years on early UNIX and other timesharing systems involved a programmer writing a fake login program. It would prompt the user for a user ID and password. No matter what the user typed, the program would indicate an invalid login attempt and then transfer control to the real login program. The spoofing program would write the login and password into a disk file, which was retrieved later.

The most popular spoofing attacks today are IP spoofing, ARP spoofing, and DNS spoofing. With IP spoofing, the goal is to make the data look as though it came from a trusted host when it didn't (thus spoofing the IP address of the sending host), as shown in Figure 5.31. The threat agent will forge their packet with the victim's source address.

With *ARP spoofing* (also known as *ARP poisoning*), the media access control (MAC) address of the data is faked. By faking this value, it is possible to make it look as though the data came from a networked device that it did not come from. This can be used to gain access to the network, to fool the router into sending to the device data that was intended for another host, or to launch a DoS attack. In all cases, the address being faked is an address of a legitimate user, making it possible to get around such measures as allow/deny lists.

With *DNS spoofing*, the DNS server is given information about a name server that it thinks is legitimate when it isn't. This can send users to a website other than the one to which they wanted to go, reroute mail, or do any other type of redirection for which data from a DNS server is used to determine a destination. Another name for this is *DNS poisoning*.

FIGURE 5.31 IP address spoofing attack

Always think of spoofing as fooling. Attackers are trying to fool the user, system, and/or host into believing that they're something they are not. Because the word *spoof* can describe any false information at any level, spoofing can occur at any level of a network.

The important point to remember is that a spoofing attack tricks something or someone into thinking that something legitimate is occurring.

On-Path Attack

Many of the attacks we're discussing can be used in conjunction with an *on-path attack*, which was previously known as a man-in-the-middle (MitM) attack. For example, the evil twin attack mentioned earlier allows the attacker to position themselves between the compromised user and the destination server. The attacker can then eavesdrop on a conversation and possibly change information contained in the conversation. Conventional on-path attacks allow the attacker to impersonate both parties involved in a network conversation. This allows the attacker to eavesdrop and manipulate the conversation without either party knowing. The attacker can then relay requests to the server as the originating host attempts to communicate on the intended path, as shown in Figure 5.32.

Password Attacks

Password attacks occur when an account is attacked repeatedly. This is accomplished by using applications known as password crackers, which send possible passwords to the account in a systematic manner. The attacks are initially carried out to gain passwords for an access or modification attack. There are several types of password attacks:

Brute-Force Attacks A *brute-force* attack is an attempt to guess passwords until a successful guess occurs. As an example of this type of attack, imagine starting to guess with


A and then going through z; when no match is found, the next guess series goes from AA to zz, and then it adds a third value (AAA to zzz). Because of the nature of this routine, this type of attack usually occurs over a long period of time. To make passwords more difficult to guess, they should be much longer than a few characters. It is recommended that you make the password at least 12 digits and complex. You should also have password lockout policies.

Dictionary Attacks A *dictionary attack* uses a dictionary of common words to attempt to find the user's password. Dictionary attacks can be automated, and several tools exist in the public domain to execute them. As an example of this type of attack, imagine guessing words and word combinations found in a standard English-language dictionary.

Rainbow Tables When passwords are stored in the operating system, they are stored in their hashed format. A hash is a one-way cryptographic algorithm. If you have access to the password hash, a rainbow table can be used to find the password. Rainbow tables are databases that have had every permutation run through a hashing algorithm, so the hash can be looked up and cross-referenced back to the password that created it. These databases can be really large. It is not uncommon for a rainbow table to ship on a portable hard drive.

Hybrid Attacks A *hybrid attack* typically uses a combination of dictionary entries and brute force. For example, if you know that there is a good likelihood that the employees of a particular company are using derivatives of the company name in their passwords, then you can seed those values into the values attempted.

FIGURE 5.32 On-path attack



Insider Threat

Insider threats are threats that originate from within your organization. Employees know the organization and can navigate the organization to get the information they need. A disgruntled employee can carry out an attack on the organization by leaking information or

selling it. When information is sold to a competitor for profit, it is considered corporate espionage. The inside threat does not always need to be criminal in intent. It can also be as simple as an employee plugging an unauthorized wireless access point into the corporate network.

SQL Injection

A *Structured Query Language (SQL) injection* attack occurs when a threat agent enters a series of escape codes along with a well-crafted SQL statement into a URL. The seamlessly harmless page on the backend that is awaiting the request runs the SQL query along with its normal query. For example, a normal post URL might look like this:

```
http://www.wiley.com/phone.php?name=jones
```

The threat agent will add their SQL injection after the normal post query string, such as the following:

```
http://www.wiley.com/phone.php?name=jones; DROP TABLE Users
```

This would generate the following SQL query on the backend and send the malicious query to the SQL database:

```
SELECT FullName, PhoneNum  
From Phones  
Where FullName Like '%jones%'; DROP TABLE USERS
```

The first two and half lines to the semicolon are generated by the page the query is posted to. The line basically tells the SQL database to return the full name and phone number for anything that contains jones. However, the threat agent appended `DROP TABLE Users` to the query with a semicolon. This will delete the Users table and cause disruption. Technically the SQL injection causes a DoS attack. Other malicious queries that are not so obvious and disruptive can be submitted to discover information like table structure and consequently steal data. Many retailers, banks, and online stores, just to name a few, have fallen prey to SQL injection attacks and made front-page news.

The best way to combat this attack is by building input validation into the rendered page on the backend; this is also known as *sanitization*. This mitigation tactic is well outside the scope of the exam, but understanding the attack is the key takeaway.

Cross-Site Scripting (XSS)

Cross-site scripting (XSS) is a tactic a threat agent uses to deliver a malicious script to the victim by embedding it into a legitimate web page. Common delivery methods for XSS are message boards, forums, or any page that allows comments to be posted. The threat agent will submit a post to these types of pages with their malicious script, such as JavaScript. When the victim browses the page, the threat agent's script will execute.

JavaScript and other scripting languages are controlled tightly by the browser, so direct access to the operating system is usually not permitted. However, the script will have access

to the web page you are browsing or the cookies the actual page stores. This type of attack is common in hijacking web pages and trying to force the user into installing a piece of malware.

Business Email Compromise (BEC)

Business email compromise (BEC) is a form of cybercrime that targets organizations by exploiting them through email. Threat actors either use compromised legitimate business email accounts, or create convincing spoofed accounts to deceive employees, partners, or clients into transferring funds or sensitive information. This type of attack relies heavily on social engineering tactics and detailed research to make fraudulent requests appear authentic. The components of a BEC attack are as follows:

Research and Targeting Threat actors begin by gathering information about the target organization. They use publicly available data from company websites, social media, and other online sources to identify key individuals, such as executives or finance employees.

Email Compromise or Spoofing The threat actor either compromises the email accounts of these key employees through phishing or other means. If the threat actor cannot compromise a legitimate account, they will create spoofed email addresses that closely resemble a legitimate account.

Crafting Deceptive Emails Using the compromised or spoofed accounts, attackers send carefully crafted emails that appear to come from trusted sources. These emails often contain urgent requests for wire transfers, sensitive data, or changes to payment details.

Exploiting Trust and Urgency The emails are designed to exploit the recipient's trust in the sender and create a sense of urgency. This makes it difficult for the recipient to recognize the fraud.

The consequences of BEC attacks can be severe financially for the organization. Financial losses can sometimes amount to millions of dollars, depending on the extent of the compromise. Additionally, the breach of trust can damage business relationships and the organization's reputation.

Supply Chain/Pipeline Attacks

A supply chain attack, also known as a pipeline attack, targets the tight integration of today's modern business operations. These attacks exploit vulnerabilities in third-party vendors or suppliers to gain access to a primary target's network. By compromising a trusted partner, the threat actor can bypass any security measures that the primary target might have in place.

One of the most infamous examples is the SolarWinds attack in 2020. Attackers inserted malicious code into the software updates for a product called Orion, which was used for IT management. When customers installed the updates, they unknowingly introduced a


backdoor into their systems. This allowed the attackers to conduct espionage and data theft on a massive scale. Many customers of the SolarWinds software were government agencies and Fortune 500 companies.

Supply chain/pipeline attacks are effective for threat actors because they can leverage the trust and access granted to third-party vendors. This makes detection and mitigation challenging, since the malicious activity often appears legitimate. To combat these threats, organizations must implement stringent security measures, such as thorough vetting of suppliers, continuous monitoring of third-party activities, and adopting a zero-trust approach to network access.

Exploits and Vulnerabilities

Exploits and vulnerabilities both have the same effect of compromising systems. Vulnerabilities are weaknesses in security for an operating system or network product. Vulnerabilities are the reason we need to constantly patch network systems. Exploits are scripts, code, applications, or techniques used in exploiting the vulnerabilities by a threat agent, as shown in Figure 5.33. In the following section we will cover the most common vulnerabilities as they pertain to the CompTIA exam.

FIGURE 5.33 Threat agents, exploits, and vulnerabilities



Noncompliant Systems

One of the easiest ways to make your systems vulnerable and expose them to threats is to fail to keep them compliant. As an administrator, you should always follow security regulatory standards as well as compliance standards.

One product that can keep your operating systems compliant is Microsoft Endpoint Configuration Manager (MECM). MECM allows for the publishing of a baseline for the Windows operating system. It will then monitor the baseline against the operating systems in your organization and will remediate them if they fall out of compliance.

MECM is just one of many tools that can be used for compliance, several of which are third-party tools. Third-party compliance solutions provide other unique benefits, such as the compliance of third-party applications in addition to Windows.

Patching and Updates

When operating systems are installed, they are usually point-in-time snapshots of the current build of the operating system. From the time of the build to the time of installation, several vulnerabilities can be published for the operating system. When an OS is installed, you should patch it before placing it into service. Patches remediate the vulnerabilities found in the OS and fixed by the vendor. Updates add new features not included with the current build. However, some vendors may include vulnerability patches in updates. Network devices also have patches and updates that should be installed prior to placing them into service.

After the initial installation of the device or operating system and the initial patches and updates are installed, you are not done! Vendors continually release patches and updates to improve security and functionality, usually every month and sometimes outside of the normal release cycle. When patches are released outside of the normal release cycle, they are called out-of-band patches and are often in response to a critical vulnerability.

Microsoft products are patched and updated through the Windows Update functionality of the operating system. However, when an administrator is required to patch and update an entire network, Windows Server Update Services (WSUS) can be implemented. A WSUS server enables the administrator to centrally manage patches and updates. The administrator can also report on which systems still need to be patched or updated.



Microsoft had officially announced the Windows Server Update Services (WSUS) is now deprecated. No new features will be developed for the product, but the feature will still function as designed. In the future this will leave a void for administrators that need internal patch management. Ultimately third-party patching products will become popular again.

When endpoints are Azure joined, Intune can provide control over update policies, allowing administrators to schedule updates, enforce compliance, and monitor the update status of all managed devices. By integrating Microsoft Intune with Windows Update for Business, organizations can take advantage of features like Delivery Optimization and Windows Update rings, which can optimize the update process and minimize disruption to end users.

Operating Systems' Life Cycle

All operating systems have a life cycle of release, support, and eventually end-of-life (EOL). When most operating systems reach their EOL, the vendor stops supplying security patches. The void of the most current patches creates a giant vulnerability for the organization, since the operating system is no longer protected from the latest vulnerabilities. To combat these vulnerabilities, keep the operating system current. This is accomplished with continual upgrades to the operating system as new versions are released.

Unprotected Systems

Obviously, an unpatched system presents a huge vulnerability in your network, but an unprotected system is just as vulnerable. A workstation or server without antivirus protection or firewall protection poses a significant risk. A workstation or server without antivirus software can contract malware and potentially infect other computers and ultimately leak data. A missing or misconfigured firewall is equally as vulnerable. The firewall is used to prevent unauthorized connections that could exploit a vulnerability in the operating system. Firewalls are also often used in place of an undesirable patch that causes other issues. For example, turning on a firewall rule that prevents connecting to the Print Spooler service on Windows is a good protection method. It can shield you from print spooler vulnerabilities if the workstation or server is not providing print services. However, if the operating system is not protected by the firewall and an exploit is released, it is consequently vulnerable.

BYOD

Security is the biggest concern as it applies to BYOD devices. The biggest reason is that the organization has less control over BYOD devices than over devices it issues and owns. BYOD devices come with two inherent risks: data leakage and data portability. Data leakage happens when a device is lost or compromised in some way. There are tactics to mitigate this, such as full device encryption. However, the user's device is then forcefully encrypted by the organization and there could be legal ramifications. Another common tactic is to use mobile device management (MDM) software that creates a partition for company data. This would allow the company to encrypt their data and not affect user data.

Data portability means that the user can cart away organizational data when they leave. Although most of the time this is not a risk, an unscrupulous salesperson may be a big risk to the organization. A line-of-business (LOB) application should be selected that displays only the data on a mobile device and does not allow data storage. Another tactic is to employ MDM software that allows remote wiping of the organization's data. When an employee leaves, the wipe is executed and the organization's data is gone. This type of functionality is also useful if a device is lost, so it also mitigates the risk of data leakage.



Exam Objective 2.7, “Given a scenario, apply workstation security options and hardening techniques,” requires you to know data encryption, password considerations, and end-user best practices and account management.

Security Best Practices

A best practice is a technique or methodology that has largely been adopted by the professional community. A number of security best practices are techniques that can keep the


organization safe and secure. Many of the best practices will mitigate some of the risk from attacks that I previously discussed. The benefits of other best practices might not be immediately apparent, but they will keep an organization safe in the long run.

Data Encryption

To prevent the loss of data, data encryption should be considered—not that the data is really ever lost, but it's no longer within your control. Consider an example of a laptop with sensitive patient record information stored on it. If the laptop were to be stolen, there are a number of utilities that could provide unauthorized access. However, with encryption (such as BitLocker) enabled, both the operating system and the data would remain encrypted and inaccessible.

Three concepts are associated with data encryption: data in use, data in transit, and data at rest, as shown in Figure 5.34. *Data in use* is the concept of data that is in an inconsistent state and/or currently resident in memory. Most of the time, you don't need to be too concerned with data in memory, since that is normally a function of the operating system. However, when data is written to a temporary location, it is considered data in use and therefore should be encrypted. *Data in transit* is information traversing the network and should always be encrypted so that it is not intercepted. Over the past decade, just about every website and application has adopted some form of encryption, so there is no reason not to use encryption in transit.

FIGURE 5.34 Data and encryption



Data at rest is a point of contention because it is believed that once the data hits the server, it's safe. However, the data is more vulnerable because it's in one spot. If a drive needed to be replaced because it went bad, outside of physical destruction there is no way to assure the data is inaccessible. When used for backup tapes, it is not only a good idea but should be a requirement.

Setting Strong Passwords

One of the most effective ways to keep a system safe is to employ strong passwords and educate your users about their best practices. Many password-based systems use a one-way hashing approach. You can't reverse the hash value in order to guess the password. This makes it impossible to reverse the hashes if the database of stored passwords is lifted (stolen) from the

operating system. Because the hash is sent over the network in lieu of the actual password, the password is harder to crack.

Passwords should be as long and complex as possible. The current recommendation by the National Institute of Standards and Technology (NIST) is a password should be at least 15 characters. Keep in mind, you should use 20 or more characters if security is a real concern. If you use only the lowercase letters of the alphabet, you have 26 characters with which to work. If you add the numeric values 0–9, you get another 10 characters. Adding uppercase letters, you gain an additional 26 characters. If you go one step further by using symbol characters (such as !"#\$%&'()*+,-./:;<=>?@[\\]^_`{|} and ~, including a blank space), you have an additional 33 more characters. You then have a pallet of 95 characters for each position in your password. A typical example of a complex password using all these elements might be %s4@7dFs#D2\$. If you have a hard time coming up with a strong password on your own, you can always use an online password generator, such as <https://passwordsgenerator.net>.

When the *Password Complexity* policy in Group Policy is enabled for the Windows operating system, three of the four categories—lowercase, uppercase, numbers, and symbols—must be used in your password. Windows Server and/or Active Directory can also require a minimum password size, which guarantees a secure password when coupled with complexity.



Most vendors recommend that you use symbols in your password, and some go so far as to require it. Some symbols, however, do not work for web-based applications. Therefore, you may find that a vendor accepts a reduced set of characters. Although this reduces the complexity, using characters in addition to letters and numbers makes the password more secure.

Let's look further at password complexity. If you used a 4-character password, this would be $95 \times 95 \times 95 \times 95$, or approximately 81.5 million possibilities. A 5-character password would give you 95^5 , or approximately 7.7 billion possibilities. And a 10-character password would give you 95^{10} , or 5.9×10^{18} (a *very* big number) possibilities. As you can see, these numbers increase exponentially with each position added to the password. A 4-digit password could probably be broken in a fraction of a day, whereas a 10-digit password would take considerably longer and much more processing power.


If your password consisted of only the 26 lowercase letters, the 4-digit password would have 26^4 , or 456,000 combinations. A 5-character password would have 26^5 , or over 11 million combinations, and a 10-character password would have 26^{10} , or 1.4×10^{14} combinations. The number of combinations is still a big number, but it would take considerably less time to break it compared to a longer password. This is all based on the notion that a brute-force password attack is being performed. If a dictionary attack were being performed, a 4- or 5-digit lowercase password could take less than 5 minutes to crack.

Mathematical methods of encryption are primarily used in conjunction with other encryption methods as part of authenticity verification. The message and the hashed value of the message can be encrypted using other processes. In this way, you know that the message is secure and hasn't been altered.

Requiring Passwords

Make absolutely certain that you require passwords for all accounts. It's such a simple thing to overlook in a small network, but it's not something a malicious user will overlook. By default, Windows will not allow an account to connect over the network if it has a blank password. It will, however, allow a person to log in locally with a blank password. There is a security option in the local Group Policy that specifies this behavior, as shown in Figure 5.35.

FIGURE 5.35 Windows Group Policy security options



The operating system is not the only place where you should use a password for security. You should also use passwords on the *basic input/output system (BIOS)* and *Unified Extensible Firmware Interface (UEFI)* firmware. If a malicious user has access, they could possibly circumvent your security by booting a live operating system.


You should also change the *default passwords* on system accounts. There are dedicated sites on the web that document the default username and password for various vendor devices. A common hacker can easily pull up these sites and find the default username and password for your wireless access point, camera system, or any other system on which you've neglected to change the default password.

Password Expiration

Password expiration should be a consideration because passwords can be compromised as time goes on. Whether passwords are compromised by shoulder surfing or keylogging, or intercepted via the network when a user is logging in, the fact remains they are only one factor of authentication. Therefore, passwords should be set to expire on a monthly, bi-monthly, quarterly, semi-annual, or annual basis. The more sensitive an account is, the more frequently the password should be changed.

Windows has a default password expiration of 42 days, as shown in Figure 5.36. You should put in place a system to expire passwords on a periodic basis, as stated previously. You would then communicate this to your users via the onboarding process when they are hired.

FIGURE 5.36 Password expiration



End-User Best Practices

In addition to administrator best practices, there are several different end-user best practices that you should advocate to your users. In the following, we will cover the top end-user best practices covered by the CompTIA exam. However, when it comes to end-user best practices and training, these are just the tip of the iceberg.

Locking Screens

When a user walks away from their computer and leaves themselves logged in, anyone who walks up to the computer has the same level of access as the owner of the account. This type of attack requires that the threat agent be physically present. However, leaving a computer logged in also invites insider threats, unauthorized access to information, or even data loss.

Training users to lock their screen when they walk away is the best way to prevent unauthorized access. By simply pressing the Windows key and L, a user can lock their screen as they walk away.

Alternately, the administrator can require a user to use a screen saver lock. For example, the screen saver lock can be set to 15 minutes. After 15 minutes of idle time, the screen saver will turn on. The user will not be able to access the desktop until they enter their password. This setting provides two benefits: First, it provides a visual deterrent to potential threat agents, and second, it prevents threat agents from carrying out an attack.

Logging Off

When users are not utilizing a system, they should be encouraged to log off the system. When users remain logged in, the programs that they were running stay running as well. If there is malware on the system, it will stay running as well, potentially allowing threat agents to carry out attacks.

When a user logs off the operating system, any malware running will terminate and hopefully not launch on next login. Malware that launches on the next login is considered to be persistent. Outside of malware, if the system has a resource that is shared, then having users log off will free the resource for the next person.

The administrator has control at their disposal that allows them to police the user. After a period of time in which the system is idle, the administrator can forcibly log off the user automatically. This is usually performed on shared systems, such as a terminal server that serves applications or virtual desktop infrastructure that serves desktops.

Securing Equipment and Information

It is our job as administrators to protect information, such as personally identifiable information (PII), as well as usernames and passwords. However, we also bestow this responsibility onto our users, since many times they have direct access to information. Users should be trained to identify PII and methods to protect such information. Examples of end-user measures to protect sensitive information can be as simple as controlling printouts, using discretion when viewing information with others around, and destroying sensitive trash, just to name a few.

End users also have portable devices that can contain sensitive data, and these devices should be secured when not in use. Many an organization has made front-page headlines with the loss of a simple laptop containing PII. Locks can be used to secure devices, and most portable devices contain a security slot for a lock. A laptop with sensitive data is not the only device that can be lost or stolen—portable hard drives can also contain PII and thus should be physically controlled, or their use should be prevented.

Password Managers

Password managers offer a solution for managing the thousands of passwords required in today's online environment. These applications securely store and encrypt credentials in a vault. The vault can only be accessed with a master password or biometric authentication. This ensures that even if the device is compromised, the stored passwords remain protected. Most password managers also have the ability to upload the vault entries to the cloud. This allows you the ability to sync vault on any computer you are using.

One of the primary benefits of using a password manager is the ability to generate strong, unique passwords for each account. This practice significantly reduces the risk of credential stuffing attacks, where threat actors use stolen passwords from one site to gain access to other accounts. By having unique passwords for each account, the impact of a single data breach is minimized.

Password managers also provide user convenience by automatically filling in login credentials on websites and applications. This not only saves time, but also reduces the likelihood of phishing attacks. Since the password manager will only autofill credentials on legitimate sites.

Another common feature of password managers is their ability to store other sensitive information, such as credit card details, secure notes, and personal identification numbers (PINs). This provides a single, secure location for all critical data.



The macOS Keychain Access app, and its successor, the Passwords app, was covered in Volume 2 Chapter 4, "Working with macOS and Linux." Although the Windows Credential Manager is not a true password manager, it stores web credentials for applications running on the Windows operating system.

Account Management

Given a security-related scenario, account management can take into consideration such settings as restricting user permissions, setting login time restrictions, disabling the Guest account, locking an account after a certain number of failed attempts, and configuring a screen lock when the system times out after a specified length of inactivity.


Restricting User Permissions

When assigning user permissions, follow the *principle of least privilege*: give users only the bare minimum that they need to do their job. Assign permissions to groups rather than to users, and make users members of groups (or remove them from groups) as they change roles or positions.

The use of groups is crucial to account management, because when you apply permission on NTFS for the user, you need to visit the resource to identify the permissions granted. When you use a group and apply the NTFS permission to that specific group, you can

then look at either the membership of the group or the membership of the user in Active Directory. This allows you to see who has access to the resource without having to visit the resource. Figure 5.37 shows an oversimplified example. A user, Fred, is a member of both the Sales and R&D groups; therefore, he has access to the Sales and R&D folders. In a real-world application, the group would be more descriptive, such as Perm_RW_Sales_Server1. This naming would describe what the group is used for (permissions), what level of permissions (RW), what resource (Sales share), and on what server (Server1).

FIGURE 5.37 Users, groups, and resources




Setting Time Restrictions

Configure user accounts so that logins can occur only during times that the user can be expected to be working. Preventing logins at 2:00 a.m. can be an effective method of keeping hackers from your systems. This can be performed in Active Directory by clicking the user's account, selecting the Account tab, and then clicking Logon Hours, as shown in Figure 5.38. From this interface, you can configure the permitted hours for logins and denied hours for logins.

Account Expiration

You can also set account expiration on the Properties tab of the user. By default, the account is set to never expire. However, you can add a date at which time the account will expire. This is best used on contractor accounts, where the terms of use can be defined. Because contractor accounts often don't go through typical human resource processes, the accounts can be forgotten about. By adding an account expiration, you can be assured that at the end of the contract the account will be disabled (expired).

FIGURE 5.38 Account restrictions

Limiting the Number of Failed Login Attempts


Configure user account settings to limit the number of login attempts before the account is locked for a period of time. Legitimate users who need to get in before the block expires can contact the administrator and explain why they weren't able to give the right password three times in a row, and illegitimate users will go away in search of another system to try to enter.

When choosing the number of failed attempts, you need to consider the number of calls you get to the help desk versus the security in having few failed attempts before lockout. You'll find that when you set it to three failed attempts, the help desk will get more calls than necessary, but it allows for better security. Setting the number of failed login attempts to five may be better for users, because many users realize after the third failed attempt that their Caps Lock key was on, but it's less secure than three failed attempts. This setting needs to be evaluated against your security requirement and help desk volume.

You should also consider the length of the lockout. If it's a Monday morning and a person enters their password wrong X number of times and gets locked out, 5 minutes

might be appropriate. The time it takes to get a cup of coffee and unlock might be just enough time on a Monday morning to allow the user to wake up. You can specify these settings for an entire domain, as shown in Figure 5.39. As shown here, the user will be locked out for 30 minutes after three failed attempts. By default, there is no account lockout policy set for a domain.

FIGURE 5.39 Account lockout policy settings



If a password attack is being executed, no protection could easily allow millions of password attempts. By setting a failed login attempt counter and a lockout duration, you slow the progress of the attack.

Changing Default Usernames

Default accounts represent a huge weakness because everyone knows they exist. When an operating system is installed—whether on a workstation or a server—certain default accounts are created. Knowing the names of those accounts simplifies the process of potential attackers accessing them because they only have to supply the password.

The default username should be changed whenever possible. Several websites are dedicated to documenting the default username and password for routers, switches, and other network equipment. These sites are useful, especially when you lose the documentation for the device and have to *factory-reset* it. A login after you have successfully just factory-reset the device is

really the only time that it's acceptable to use the default username and password for security. You don't need a website to guess the common administrative accounts on equipment or an operating system. They are usually admin, administrator, root, or sysadmin. Changing the default username makes it more challenging for someone to try to guess the credentials.


Changing the default password to a complex password is also a good practice in hardening the device. However, changing the username will also ensure that a brute-force attack cannot be performed against the default username. There are many different websites dedicated to listing the default credentials for network devices, so it doesn't take tremendous skill to obtain the default username and password of a device.

Disabling the Guest Account

When Windows is installed, one of the default accounts it creates is Guest. This represents a weakness that can be exploited by an attacker. While the account cannot do much, it can provide initial access to a system, which the attacker can use to find another account or acquire sensitive information about the system.

You should disable all accounts that are not needed, especially the Guest account. Starting with Windows 10, the operating system disables the Guest account by default, as shown in Figure 5.40. After you disable accounts that are not needed, rename the accounts, if you can. (Microsoft won't allow you to rename some.) Finally, change the passwords from the defaults and add them to the list of passwords that routinely get changed.

FIGURE 5.40 Disabled Guest account




Requiring Screen Saver Passwords

You can require lengthy, complex passwords for all your users as well as lock down the operating system with passwords, but if a user goes to the restroom or on break without locking their workstation, any person wandering by can access everything the user has privileges to.

A *screen saver* should automatically start after a short period of idle time, and a password should be required before the user can begin the session again. This method of locking the workstation adds one more level of security. A Group Policy can be put in place to turn on password-protected screen savers. Adding a password-protected screen saver can ensure that if a workstation is left unattended, it will lock and require a password to resume access. You can access this setting on Windows 11 by right-clicking an empty portion of the Desktop and selecting Personalize > Lock Screen > Screen Saver. Then, in the Screen Saver Settings dialog box, to manually require a password after a screen saver has activated, select the On Resume, Display Logon Screen check box, as shown in Figure 5.41.

FIGURE 5.41 The Screen Saver Settings dialog box



Disable AutoRun

It is never a good idea to put any media in a workstation if you do not know where it came from or what it is. The simple reason is that said media (CD, DVD, USB) could contain malware. This attack is commonly referred to as a *drop attack*. Compounding matters, that malware could be referenced in the `autorun.inf` file, causing it to be summoned when the

media is inserted in the machine and requiring no other action. `autorun.inf` can be used to start an executable, access a website, or do any of a large number of different tasks. The best way to prevent a user from falling victim to such a ploy is to disable the *AutoRun* feature on the workstation.

Microsoft has changed the function on Windows so that it no longer acts as it previously did (prior to Windows 7). The feature is now disabled by default. The reason Microsoft changed the default action can be summed up in a single word: security. That text-based `autorun.inf` file can not only take your browser to a web page, it can also call any executable file, pass along variable information about the user, or do just about anything else imaginable. Simply put, it is never a good idea to plug any media into your system if you have no idea where it came from or what it holds. Such an action opens up the user's system—and the network—to any number of possible risks. An entire business's data could be jeopardized by someone with elevated privileges inadvertently putting a harmful USB thumb drive or CD into a computer at work.

AutoPlay

The AutoRun feature is disabled by default, so that malicious software does not start automatically. However, the ability to automatically start a video, music, or open a folder when removable media is inserted into the computer is really useful. The functionality of automatically performing an action or asking what should be done when media is inserted is still enabled by default on Windows, via a feature called *AutoPlay*. AutoPlay distinguishes itself from AutoRun, since it does not look at the `autorun.inf` file and does not start an executable unless the user specifically clicks an option that they want to do so.



Exam Objective 2.9, “Compare and contrast common data destruction and disposal methods,” requires you to know the various methods for destroying hard drives. In addition, you should know the recycling and repurposing best practices, the outsourcing of disposal, as well as regulatory and environmental requirements.

Destruction and Disposal Methods

Think of all the sensitive data written to a hard drive. The drive can contain information about students, clients, users—about anyone and anything. The hard drive can be in a desktop PC, in a laptop, or even in a printer. Many laser printers above consumer grade offer the ability to add a hard drive to store print jobs. If the drive falls into the wrong hands, you can not only lose valuable data but also risk a lawsuit for not properly protecting privacy. An appropriate data destruction/disposal plan should be in place to avoid any potential problems.

Since data on media holds great value and liability, that media should never simply be tossed away for prying eyes to stumble on. For the purpose of this objective, the media in question is hard drives, and there are three key concepts to understand with regard to

them: formatting, sanitation, and destruction. Formatting prepares the drive to hold new information (which can include copying over data already there). Sanitation involves wiping the data off the drive, whereas destruction renders the drive no longer usable.



Although this objective is heavily focused on hard drives, it is also possible to have data stored on portable flash drives, backup tapes, CDs, or DVDs. In the interest of security, it is recommended that you destroy them before disposing of them.

Recycling or Repurposing Best Practices

For exam purposes, the best practices for recycling or repurposing fall into the categories of low-level formats (as opposed to standard formatting), overwrites, and drive wipes.

Low-Level Format vs. Standard Format

There are multiple levels of formatting that can be done on a drive. A standard format, accomplished using the operating system's format utility (or similar), can mark space occupied by files as available for new files without truly deleting what was there. Such erasing—if you want to call it that—doesn't guarantee that the information isn't still on the disk and recoverable.

A low-level format (typically accomplished only in the factory) can be performed on the system, or a utility can be used to completely wipe the disk clean. This process helps to ensure that information doesn't fall into the wrong hands.

The manufacturer performs a low-level format on hard drives. Low-level formatting must be performed even before a drive can be partitioned. In low-level formatting, the drive controller chip and the drive meet for the very first time and learn to work together. Because controllers are integrated into SATA and IDE drives, low-level formatting is a factory process. Low-level formatting is not operating system-dependent.



Never perform a low-level format on SATA, or SCSI drives! These drives are formatted at the factory, and you may cause problems by using low-level utilities on them.

The main thing to remember for the exam is that most forms of formatting included with the operating system do not actually erase the data completely. Formatting the drive and then disposing of it has caused many companies problems when individuals who never should have seen it retrieve the data using applications that are commercially available.

Hard Drive Sanitation and Sanitation Methods

A number of vendors offer hard drives with *Advanced Encryption Standard (AES)* cryptography built in. However, it's still better to keep these secure hard drives completely out of the

hands of others than to trust their internal security mechanisms once their usable life span has passed for the client. Some vendors include freeware utilities to erase the hard drive. If it is a Serial ATA (SATA) drive, you can always run HDDEraser, but you are still taking your chances.

Sanitation Utilities

HDDEraser is a freeware utility that is included with many different boot images. HDDEraser can be downloaded with Hiren's BootCD from www.hirensbootcd.org.

If you were to perform a web search, you would find several other sanitization utilities like HDDEraser, but it is important to recognize and acknowledge that many of these do not meet military or GSA specifications. Those specifications should be considered as guidelines to which you must also adhere when dealing with your own, or a client's, data. The only surefire method of rendering the hard drive contents completely eradicated is physical destruction.

Solid-state drives (SSDs) pose a greater problem since the media is flash memory and not mechanical, like conventional hard disk drives (HDDs). Low-level formats can be performed, as mentioned in the preceding section, but the 1s and 0s will still be technically on the flash memory. Therefore, many vendors have a sanitization utility for scrubbing information from SSDs. It is best to check with the vendor, as these tools are specific to the vendor and model of SSD.

Overwrite

Overwriting the drive entails copying over the data with new data. A common practice is to replace the data with 0s. A number of applications allow you to recover what was there prior to the last write operation, and for that reason, most overwrite software will write the same sequence and save it multiple times.

DBAN is a utility that comes with its own boot disk from <https://dban.org>. You can find a number of other software “shredders” by doing a quick web search.

Drive Wipe

If it's possible to verify beyond a reasonable doubt that a piece of hardware that's no longer being used doesn't contain any data of a sensitive or proprietary nature, then that hardware can be recycled (sold to employees, sold to a third party, donated to a school, and so on). That level of assurance can come from wiping a hard drive or using specialized utilities.

If you can't be assured that the hardware in question doesn't contain important data, then the hardware should be destroyed. You cannot, and should not, take a risk that the data your company depends on could fall into the wrong hands.

Physical Destruction

Physically destroying the drive involves rendering it no longer usable. While the focus is on hard drives, you can also physically destroy other forms of media, such as flash drives and optical media.

Shredder

Many commercial paper shredders are also capable of destroying DVDs and CDs. Paper shredders, however, are not able to handle hard drives; you need a shredder created for just such a purpose. A low-volume hard drive shredder that will destroy eight drives per minute can carry a suggested list price of around \$20,000.

Drill/Hammer

If you don't have the budget for a hard drive shredder, you can accomplish similar results in a much more time-consuming way with a power drill. The goal is to physically destroy the platters in the drive. Start the process by removing the cover from the drive—this is normally done with a Torx driver. (Although #8 does not work with all drives, it is a good one to try first.) You can remove the arm with a slotted screwdriver and then the cover over the platters using a Torx driver. Don't worry about damaging or scratching anything—nothing is intended to be saved. Everything but the platters can be tossed away.

As an optional step, you can completely remove the tracks using a belt sander, grinder, or palm sander. The goal is to turn the shiny surface into fine powder. Again, this step is optional, but it adds one more layer of assurance that nothing usable remains. Always wear eye protection and be careful not to breathe in any of the fine particles that are generated during the grinding/destruction process.

Following this, use the power drill to create as small a set of particles as possible. A drill press works much better for this task than trying to hold the drive and drill it with a handheld model.

Do You Really Want to Do It Yourself?

Even with practice, you will find that manually destroying a hard drive is time consuming. There are companies that specialize in this and can do it efficiently. One such company is Shred-it, which will pick up your hard drive and provide a chain-of-custody assurance and a certificate of destruction on completion. You can find out more about what they offer at www.shredit.com.

Electromagnet (Degaussing)

A large electromagnet can be used to destroy any magnetic media, such as a hard drive or backup tape set. The most common of these is the degaussing tool. *Degaussing* involves

applying a strong magnetic field to initialize the media. (This is also referred to as *disk wiping*.) This process helps ensure that information doesn't fall into the wrong hands.

Degaussing involves using a specifically designed electromagnet to eliminate all data on the drive, including the factory-prerecorded servo tracks. You can find wand model degaussers priced at just over \$500 or desktop units that sell for up to \$30,000.

Incineration

A form of destruction not to be overlooked is incineration. It is possible to destroy most devices by burning them up, using an accelerant such as gasoline or lighter fluid to aid the process. However, burning hard drives at extremely high temperatures, typically between 870°C and 1,200°C causes unrecoverable data. The intense heat causes the materials inside the hard drive, such as aluminum or glass platters, to melt and deform, ensuring that the data cannot be recovered.

This method is typically used for securely disposing of sensitive information. However, there are environmental considerations and regulations that must be followed. Many professional data destruction services offer incineration as a secure disposal option.



Be careful with any fire, particularly those in which accelerants are used. Be sure that you are not burning anything capable of releasing toxic fumes and that you have the fire controlled and contained at all times.

Environmental and Regulatory Concerns

The proper disposal of e-waste is crucial for both environmental and regulatory reasons. Environmentally, the improper disposal of computer equipment can lead to the release of hazardous materials like lead, mercury, and cadmium. These materials can contaminate soil and water, posing significant health risks to humans and wildlife. Additionally, e-waste contains valuable materials that can be recycled and reused, reducing the need for new raw materials. From a regulatory perspective, many governments have strict laws governing the disposal of e-waste to prevent environmental damage. Compliance with these regulations not only helps protect the environment but also shields organizations from potential legal penalties and fines.

Certificate of Destruction

A *certificate of destruction* (or *certificate of recycling*) may be required for audit purposes. Such a certificate, usually issued by the organization carrying out the destruction, is intended to verify that the asset was properly destroyed and usually includes serial numbers, type of destruction done, and so on.

Securing Physical Documents/Passwords/Shredding

The type and amount of information that can be gleaned from physical documents is amazing, even in the age when there is such a push to go paperless. *Dumpster diving* is a common problem that puts systems at risk. Companies normally generate a huge amount of paper,

most of which eventually winds up in dumpsters or recycle bins. Dumpsters can contain highly sensitive information (such as a password a user has written on a piece of paper because they haven't memorized it yet).

In high-security and government environments, sensitive papers should either be shredded or burned. Most businesses don't do this. In addition, the advent of "green" companies has created an increase in the amount of recycled paper, which can often contain all kinds of juicy information about a company and its individual employees.



Real World Scenario

Document-Shredding Services

Many small companies purchase document-shredding equipment and perform this task themselves. Larger companies with a higher volume can contract a shredding service. These services place document receptacles in the offices for workers to dispose of sensitive documents. The service will then come in and dispose of the documents in a controlled fashion, sometimes shredding the documents on the premises with special equipment. These services can sometimes furnish certifications of disposal, depending on the type of destruction you require.

Summary

In this chapter, you learned about the various issues related to security that appear on the A+ 220–1202 exam. Security is a popular topic in computing, and the ways in which a troublemaker can cause harm increase regularly. CompTIA expects everyone who is A+ certified to understand the basic principles of security and be familiar with solutions that exist.

You also learned of security problem areas and issues that can be easily identified. Problem areas include viruses, Trojans, worms, and malware. Security solutions include implementing encryption technology, using authentication, implementing firewalls, and incorporating security at many levels.

Exam Essentials

Know the various security measures and their purpose. There are several different physical security measures that can be applied to the public and your staff. In addition, there are several different logical security measures, such as access control lists (ACLs), multifactor

authentication (MFA), and several others discussed in this chapter. Be familiar with their practical application.

Know the various physical security devices. You should be familiar with the various physical security devices, such as the use of an access control vestibule, badge readers, video surveillance, alarm systems, motion sensors, door locks, equipment locks, guards, bollards, and fences. In addition, you should be able to describe security devices as they relate to staff, such as the use of key fobs, smartcards, keys, biometrics, lighting, and magnetometers.

Know the various logical security methods. You should be familiar with logical methods of security, such as the implementation and use of access control lists (ACLs), implementing principle of least privilege, and multifactor authentication (MFA), including email, hard tokens, soft tokens, Short Message Service (SMS), voice calls, and authenticator applications as factors of authentication.

Be able to describe why antivirus/antimalware software is needed. Antivirus/antimalware software looks at a virus and takes action to neutralize it based on a virus definition database. Virus definition database files are regularly made available on vendor sites.

Understand the need for user education. Users are the first line of defense against most threats, whether physical or digital. They should be trained on the importance of security and how to help enforce it.

Know the types of malware and mitigation. You should be familiar with the various types of malware commonly found, such as Trojan, rootkit, virus, spyware, ransomware, keylogger, boot sector virus, and cryptominers. In addition, you should be familiar with the various ways to mitigate malware with the use of the recovery console, antivirus, anti-malware, software firewalls, antiphishing training, user education, and OS reinstallation.

Know the various best practices for security. You should know the best security practices, such as implementing data-at-rest encryption, password complexity, and password expiration, as well as protecting BIOS or UEFI with passwords. In addition, you should know the various end-user best practices, such as using screen saver locks, logging off when a computer or application is not in use, and implementing security and protection of critical hardware as well as security of personally identifiable information (PII) and passwords.

Know how to compare and contrast social engineering, threats, and vulnerabilities. Social engineering variants include shoulder surfing (watching someone work) and phishing (tricking someone into believing they are communicating with a party other than the one with whom they are communicating). Variations on phishing include vishing, shoulder surfing, whaling, tailgating, impersonation, dumpster diving, and evil twin.

Know the various threats to a network. You should be familiar with the various threats to a network, such as distributed denial-of-service (DDoS), denial-of-service (DoS), zero-day attack, spoofing, on-path attacks, brute-force password attacks, dictionary password attacks, insider threats, Structured Query Language (SQL) injection, cross-site scripting (XSS), Business email compromise (BEC), and supply chain/pipeline attack.

Understand the need for good passwords. Passwords are the first line of defense for protecting an account. A password should be required for every account, and strong passwords should be enforced. Users need to understand the basics of password security and work to keep their accounts protected by following company policies regarding passwords.

Understand the difference between standard and low-level formatting. Standard formatting uses operating system tools and marks the drive as available for holding data without truly removing what was on the drive (thus, the data can be recovered). A low-level format is operating system-independent and destroys any data that was on the drive.

Understand how to implement appropriate data destruction and disposal methods. A hard drive can be destroyed by tossing it into a shredder designed for such a purpose, or it can be destroyed with an electromagnet in a process known as degaussing. You can also disassemble the drive and destroy the platters with a drill or other tool that renders the data irretrievable. Third party vendors can also be used and they will supply certification of destruction/recycling. Proper disposal is required by regulatory and environmental requirements.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which component of physical security addresses outer-level access control?
 - A. Fences
 - B. Access control vestibule
 - C. Multifactor authentication
 - D. Strong passwords
2. Which type of device can detect weapons on a person entering a facility?
 - A. Biometrics
 - B. Magnetometer
 - C. Motion sensor
 - D. Badge reader
3. As part of your training program, you're trying to educate users on the importance of security. You explain to them that not every attack depends on implementing advanced technological methods. Some attacks, you explain, take advantage of human shortcomings to gain access that should otherwise be denied. Which term do you use to describe attacks of this type?
 - A. Social engineering
 - B. IDS
 - C. Perimeter security
 - D. Biometrics
4. You're in the process of securing the IT infrastructure by adding fingerprint scanners to your existing authentication methods. This type of security is an example of which of the following?
 - A. Access control
 - B. Physical barriers
 - C. Biometrics
 - D. Softening
5. Which type of attack denies authorized users access to network resources?
 - A. DoS
 - B. Worm
 - C. Trojan
 - D. Social engineering

6. As the security administrator for your organization, you must be aware of all types of attacks that can occur and plan for them. Which type of attack uses more than one computer to attack the victim?
 - A. DoS
 - B. DDoS
 - C. Worm
 - D. Rootkit
7. A vice president of your company calls a meeting with the IT department after a recent trip to competitors' sites. She reports that many of the companies she visited granted access to the operating system or applications after an employee presented a number that rotated. Of the following, which technology relies on a rotating number for users for authentication?
 - A. Smartcard
 - B. Biometrics
 - C. Geofencing
 - D. Token
8. You've discovered that credentials to a specific application have been stolen. The application is accessed from only one computer on the network. Which type of attack is this most likely to be?
 - A. On-path attack
 - B. Zero-day
 - C. Denial-of-service (DoS)
 - D. Smurf
9. A junior administrator comes to you in a panic. After looking at the log files, he has become convinced that an attacker is attempting to use a legitimate IP address to disrupt access elsewhere on the network. Which type of attack is this?
 - A. Spoofing
 - B. Social engineering
 - C. Worm
 - D. Password
10. Which of the following is different from a virus in that it can reproduce itself, is self-contained, and doesn't need a host application to be transported?
 - A. Worm
 - B. Smurf
 - C. Phish
 - D. Trojan

11. A reflective attack attempts to use a broadcast ping on a network. The return address of the ping may be that of a valid system in your network. Which protocol does the reflective attack use to conduct the attack?
 - A. TCP
 - B. IP
 - C. UDP
 - D. ICMP
12. Which type of attack involves passing a database query with a web request?
 - A. Insider threat
 - B. Evil twin
 - C. SQL injection
 - D. Tailgating
13. Which is an example of an authentication method in which you have something?
 - A. Password
 - B. Key fob
 - C. Fingerprint
 - D. Place
14. You need to protect your users from potentially being phished via email. Which of the following should you use to protect them?
 - A. Antivirus software
 - B. End-user education
 - C. SecureDNS
 - D. The principle of least privilege
15. Your help desk has informed you that they received an urgent call from the vice president last night requesting his login ID and password. When you talk with the VP today, he says he never made that call. What type of attack is this?
 - A. Spoofing
 - B. Replay
 - C. Social engineering
 - D. Trojan horse
16. Internal users suspect there have been repeated attempts to infect their systems, as reported to them by pop-up messages from their antivirus software. According to the messages, the virus seems to be the same in every case. What is the most likely culprit?
 - A. A server is acting as a carrier for a virus.
 - B. A password attack is being carried out.
 - C. Your antivirus software has malfunctioned.
 - D. A DoS attack is under way.

17. You're working late one night and notice that the hard drive on your new computer is very active even though you aren't doing anything on the computer and it isn't connected to the Internet. What is the most likely suspect?
- A. A spear phishing attack is being performed.
 - B. A virus is spreading in your system.
 - C. Your system is under a DoS attack.
 - D. TCP/IP hijacking is being attempted.
18. You're the administrator for a large bottling company. At the end of each month, you routinely view all logs and look for discrepancies. This month, your email system error log reports a large number of unsuccessful attempts to log in. It's apparent that the email server is being targeted. Which type of attack is most likely occurring?
- A. Brute-force
 - B. Backdoor
 - C. Worm
 - D. TCP/IP hijacking
19. Your boss needs you to present to upper management the need for a firewall for the network. What is the thesis of your presentation?
- A. The isolation of one network from another
 - B. The scanning of all packets for viruses
 - C. Preventing password attacks
 - D. The hardening of physical security
20. Which Active Directory component maps printers and drives during login?
- A. Home folders
 - B. Organizational unit
 - C. Login script
 - D. Microsoft Management Console (MMC)

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Calculate the complexity of a simple 8-character alphanumeric password versus a 25-character alphanumeric password with symbols.



Chapter 6

Securing Operating Systems

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **2.0 Security**
- ✓ **2.2 Given a scenario, configure and apply basic Microsoft Windows OS security settings.**
 - Defender Antivirus
 - Activate/deactivate
 - Updated definitions
 - Firewall
 - Activate/deactivate
 - Port security
 - Application security
 - Users and groups
 - Local vs. Microsoft account
 - Standard account
 - Administrator
 - Guest user
 - Power user
 - Log-in OS options
 - Username and password
 - Personal identification number (PIN)
 - Fingerprint
 - Facial recognition
 - SSO
 - Passwordless/Windows Hello



- NTFS vs. share permissions
 - File and folder attributes
 - Inheritance
- Run as administrator vs. standard user
- User Account Control (UAC)
- BitLocker
- BitLocker-To-Go
- Encrypting File System (EFS)

✓ **2.3 Compare and contrast wireless security protocols and authentication methods.**

- Protocols and encryption
 - Wi-Fi Protected Access 2 (WPA2)
 - WPA3
 - Temporal Key Integrity Protocol (TKIP)
 - Advanced Encryption Standard (AES)
- Authentication
 - Remote Authentication Dial-in User Service (RADIUS)
 - Terminal Access Controller Access-Control System (TACACS+)
 - Kerberos
 - Multifactor

✓ **2.8 Given a scenario, apply common methods for securing mobile devices.**

- Hardening techniques
 - Device encryption
 - Screen locks
 - Facial recognition
 - PIN codes
 - Fingerprint
 - Pattern



- Swipe
- Configuration profiles
- Patch management
 - OS updates
 - Application updates
- Endpoint security software
 - Antivirus
 - Anti-malware
 - Content filtering
- Locator applications
- Remote wipes
- Remote backup applications
- Failed log-in attempts restrictions
- Policies and procedures
 - MDM
 - BYOD vs. corporate-owned devices
 - Profile security requirements

✓ **2.10 Given a scenario, apply security settings on SOHO wireless and wired networks.**

- Router settings
 - Change default passwords
 - IP filtering
 - Firmware updates
 - Content filtering
 - Physical placement/secure locations
 - Universal Plug and Play (UPnP)
 - Screened subnet
 - Configure secure management access
- Wireless specific
 - Changing the service set identifier (SSID)
 - Disabling SSID broadcast



- Encryption settings
- Configuring guest access
- Firewall settings
 - Disabling unused ports
 - Port forwarding/mapping

✓ **2.11 Given a scenario, configure relevant security settings in a browser.**

- Browser download/installation
 - Trusted sources
 - Hashing
 - Untrusted sources
- Browser patching
- Extensions and plug-ins
 - Trusted sources
 - Untrusted sources
- Password managers
- Secure connections/sites – valid certificates
- Settings
 - Pop-up blocker
 - Clearing browsing data
 - Clearing cache
 - Private-browsing mode
 - Sign-in/browser data synchronization
 - Ad blockers
 - Proxy
 - Secure DNS
- Browser feature management
 - Enable/disable
 - Plug-ins
 - Extensions
 - Features



This chapter is the second of two chapters that focus primarily on security. Volume 2 Chapter 5, “Security Concepts,” covered myriad security concepts, ranging from physical security to the proper destruction of data storage devices in your organization. In this chapter, we will focus on operating system security and mobile security.

Many organizations are adopting a cloud-first initiative for their line-of-business applications. This has further perpetuated the adoption of mobile devices in the workplace. These initiatives, along with the rapid adoption of mobile devices in our personal lives, has created a tremendous need for security. This chapter will address the concerns of operating system security, mobile device security, and best practices.

Working with Windows OS Security Settings

Every operating system offers security features and settings. While you need to know a little about Linux and macOS, the A+ exams focus primarily on Windows and the OS-specific security settings that you need to know to secure them. The following sections will explore some basic Windows OS security features and settings in more detail.



Exam Objective 2.2, “Given a scenario, configure and apply basic Microsoft Windows OS security settings,” requires you to know users and groups in the Windows operating system. In addition, you should be familiar with log-in options, NTFS and share permissions, UAC, antivirus, and encryption in the Windows operating system.

Users and Groups


A number of groups are created on the operating system by default. The following sections look at the main ones.

Microsoft vs. Local Accounts

During the initial setup of Windows, Microsoft urges the user to log in with a Microsoft account, as shown in Figure 6.1. You can set up a Microsoft account with an email, a phone,

or a Skype login, which is actually a Microsoft account. This feature was originally introduced with Windows 8.

FIGURE 6.1 Microsoft account screen



Setting up your initial user account is a smart idea, because it allows you to set up tools, such as OneDrive, to back up your files. The most significant feature is it allows setup of synchronization for settings across all of your devices. However, it also allows access to other Microsoft productivity tools and ultimately it is how you access the Microsoft Store and the Store App ecosystem.

You can also select Sign In Options and choose to use an offline account, as shown under the logon name in Figure 6.1. When you use an offline account, many of the features we described earlier need to be set up manually. Also, depending on the feature, it may not work as designed. When you elect to create an offline account, you are actually electing to use local accounts on the operating system. Many corporate-owned devices still use local accounts and connect to a traditional Active Directory (AD) domain.

Choosing to sign in with a Microsoft account or local (offline) account depends on what you are trying to achieve. If the device will be used for personal work, then a Microsoft account is the best option. If the device will be used for an organization, then a local account

might be the best option. A third option exists that is reserved for organizations, which is the use of a corporately owned email address for enrollment into mobile device management (MDM) software, such as Intune.

In the following sections we will cover the basic local accounts and local permissions that you should expect to see on the CompTIA exam.

Administrator

The *Administrator* account is the most powerful of all; it has the power to do everything from the smallest task all the way to removing the operating system. Because of the great power the Administrator account holds, and the fact that it is always created, many who want to do harm target this account as the one that they try to breach. To increase security, during the installation of the Windows operating systems in question, you are prompted for the name of a user who will be designated as the Administrator. The power then comes not from being called “Administrator” (the username might now be “wiley,” “sybex,” or something similar) but from being a member of the *Administrators group*. (Notice the plural for the group and singular for the user.)

Since members of the Administrators group have such power, they can inadvertently do harm (such as accidentally deleting a file that a regular user could not). To protect against this, the practice of logging in with an Administrators group account for daily interaction is strongly discouraged. Instead, we suggest that system administrators log in with a user account (lesser privileges) and change to the Administrators group account (elevated privileges) only when necessary.

Power User

Originally, Microsoft wanted to create a group in Windows whose members were not as powerful as members of the Administrators group, so they created the *Power Users* group. The idea was that members of this group would be given Read and Write permission to the system, allowing them to install most software but keeping them from changing key operating system files. As such, it would be a good group for those who need to test software (such as programmers) and junior administrators.

The group did not work out as planned, and in Windows 7 up to Windows 11 the group has no more permissions than a standard user. The group is now kept around only for backward compatibility with Windows XP systems.

Guest

The *Guest user* account is created by default (and should be disabled) and is a member of the Guests group. For the most part, members of the Guests group have the same rights as a standard user, except they can’t get to log files. The best reason to make users members of the Guests group is that they can access the system only for a limited time.



As part of operating system security, it’s usually recommended that you rename the default Administrator and Guest accounts that are created at installation.

Standard User

By default, *standard users* belong to the local Users group. Members of this group have Read and Write permission to their own profile. They cannot modify systemwide Registry settings or do much harm outside of their own accounts. Under the principle of least privilege, users should be made members of the Users group only, unless qualifying circumstances force them to have higher privileges.

If you attempt to run some utilities (such as `sfc.exe`) from a standard command prompt, you will be told that you must be an administrator running a console session in order to continue. If your account is in the Administrators group, then the command prompt must be launched with the elevated permissions of the administrator. To do so, choose Start > All > Windows Tools, and then right-click Command Prompt and choose Run As Administrator. The User Account Control (UAC) will prompt you to continue, and then you can run `sfc.exe` without a problem.

Changing Between Account Types

You can change between account types of standard and administrator for your local user account or for other local user accounts by using the legacy Control Panel applet for User Accounts. You simply open Control Panel, select User Accounts, click Manage User Account, click Properties, and click the Group Membership tab to change your level of access. You can also select another user and then change their group membership. You will be presented with the dialog box in Figure 6.2, where you can change the type of account from Administrator to Standard using the radio buttons.




If the Standard option is grayed out and you cannot change the user account to a standard user, then you must make another administrative user. The operating system must have at least one valid administrative account. The operating system will not let you lock yourself out through this set of screens.


User Authentication

Users can log into the local operating system with their username and password, if they have an account, and they will receive a local access token. The *access token* the user is granted is locally significant for the operating system. For example, an administrator (local) who authenticates against the operating system is only an administrator of that operating system and has no further network permissions. Every Windows operating system has a local database and authentication system called the *Security Account Manager (SAM)*, as shown in Figure 6.3.

Active Directory simplifies the sign-on process for users and lowers the support requirements for administrators. Access can be established through groups and enforced through group memberships: all users log into the Windows domain using their centrally created Active Directory account. It's important to enforce password changes and make certain that passwords are updated throughout the organization on a frequent basis.

FIGURE 6.2 Changing the account type


Active Directory uses Kerberos v5. A server that runs Active Directory retains information about all access rights for all users and groups in the network. When a user logs into Active Directory, they are granted a *network access token*, also called a *Kerberos token*. This token can be used to authenticate against other servers and workstations in the domain and is accepted network (domain) wide. This token is also referred to as the user's *globally unique identifier (GUID)*. Applications that support Active Directory for authentication can use this GUID to provide access control.

FIGURE 6.3 Windows authentication

Single Sign-On (SSO)

One of the big problems larger networks must deal with is the need for users to access multiple systems or applications. This may require users to remember multiple accounts and passwords. An alternative to this is that the application must support Active Directory authentication, but that creates other considerations.

The purpose of *single sign-on (SSO)* is to give users access to all the applications and systems that they need when they log in. Single sign-on is often used with cloud-based resources. The principle behind SSO is that the resource will trust that the user has already been authenticated. This allows the user to use one set of credentials for all resources they have access to. The authentication server performs this by sending a claim on behalf of the user, as shown in Figure 6.4. This claim can contain any number of Active Directory user attributes, such as first name, last name, email, and username, just to name a few. It is important to understand that at no time during the authentication process are the username and password sent to the resource that is requesting authentication. The resource must trust that the user has already been authenticated and accept the claim at face value.

FIGURE 6.4 Claims-based authentication

SSO is both a blessing and a curse. It's a blessing in that once users have been authenticated, they can access all the resources on the network and browse multiple folders. Single sign-on is a curse in that it removes the doors that otherwise exist between the user and various resources.

Windows Hello

The problem with traditional usernames and passwords is that they are too complex for the user to consistently type. As administrators we want to make sure that users lock their workstation and type their password when they return to gain access. However, this is a burden on the user, especially if we are constantly changing passwords and making them more complex.

Windows Hello addresses these problems by storing the user's credentials in a secure container called the Credential Manager. The Credential Manager is then locked and unlocked with the authentication of biometrics or a PIN by the user. Once the Credential Manager is unlocked, the credentials can be passed to the operating system to provide the login credentials. Windows Hello can be configured by navigating to Start > Settings > Accounts > Sign-in Options, and you will be brought to the various ways to configure Windows Hello, as shown in Figure 6.5.

There are several different sign-in options for Windows Hello that will allow for easier login of the device. The available options will depend on your hardware connected to the device. As an example, if you want to use your fingerprint, you will need a fingerprint reader. There are several options as shown in Figure 6.5, but they all basically fall into a few categories:


Biometric Sign-In Using the Facial Recognition (Windows Hello) option will allow you to use your face to unlock the operating system. If your device is equipped with a camera, it will sense you entering the scene and verify your face to unlock the device. You can also use your fingerprint to provide authentication, but you will need a fingerprint reader as previously mentioned. Both of these biometric options authenticate you based on something you are. Because it is coupled with something that you have (the device), the process actually provides two-factor authentication/multifactor authentication (2FA/MFA).

PIN The most common option for devices without a camera or fingerprint reader is to use a personal identification number (PIN). The PIN can be a simple password containing numbers such as a PIN you use for your bank. Although you can make this as complex as you wish, its intended purpose is to make it easier to log in for the end user. When you use this method of authentication, it is something that you know. By knowing the PIN and having physical access to the device, the process provides 2FA/MFA.

Security Key Using a hardware security key or token is another way of securing Windows Hello. You can configure Windows to log you in after you present a security key, which is a small USB device that is plugged into the device. This is an authentication method of something that you have. By having both the security key and physical access to the device, the process provides 2FA/MFA.

Picture Password A picture password is similar to a gesture-based password. You can upload a picture, then select points in the picture you must click on to allow authentication. This type of authentication is based on something that you know, and when combined with the physical access to the device, the process provides 2FA/MFA.

FIGURE 6.5 Windows Hello Sign-in options




In addition to local logins, Windows Hello can be used to authenticate users for a Microsoft account, Active Directory account, and Azure AD account. It can even be used to authenticate users for identity provider services, which is another way of addressing SSO.

Administrator vs. Standard User

Administrators have rights in the operating system that allows them to change the operating system. Standard users do not have these rights and can only make changes to their environment. Because the administrative rights can cause harm to the operating system or introduce security issues, care should be taken in granting your end users administrative permissions. You should exercise the principle of least privilege with granting end-user rights. This means that unless someone needs administrative privileges, they should always log in with a standard user account. This protects the local operating system and potentially the network (if domain authenticated) from the security threats covered in Volume 2 Chapter 5.

User Account Control (UAC) is a feature that was introduced in Windows Vista. It supports the principle of least privilege by logging an administrator in with minimal permissions. This is extremely handy for users who occasionally need administrative rights, such as your home operating system. For example, if extra privileges were required to modify the operating system, a prompt asking if the user wants to continue would be displayed, as shown in Figure 6.6. If the user answers Yes, then the user will receive the administrative token to complete the task. UAC allows the user to run as a standard user with the ability to escalate privileges.

FIGURE 6.6 UAC prompt



In Exercise 6.1 you will examine the security token for your user account. The exercise assumes that you are the administrator of the operating system.

EXERCISE 6.1**Examining a Security Token**

This exercise assumes that you have not turned off the UAC and that you are the administrator of the operating system.

1. Click the Start menu, type **cmd**, and then press Enter.
2. In the command prompt, type **whoami /ALL**, and then press Enter.
3. Record the privileges that you see in the output.
4. Click the Start menu and type **cmd**.
5. Right-click Command Prompt and select Run As Administrator.
6. Answer Yes to the UAC prompt.
7. In the command prompt, type **whoami /ALL**, and then press Enter.
8. Compare the privileges that you see in the output.

When you opened the first command prompt, it was opened with the standard user credentials of your account. When you opened the second command prompt, you selected Run As Administrator. The second command prompt was opened with the administrative security token.

NTFS vs. Share Permissions

The *New Technology File System (NTFS)* was introduced with Windows NT to address security problems. Before Windows NT was released, it had become apparent to Microsoft that a new filesystem was necessary to handle growing disk sizes, security concerns, and the need for more stability. NTFS was created to address those issues.

Although the *File Allocation Table (FAT)* filesystem was relatively stable if the systems that were controlling it kept running, it didn't do well when the power went out or the system crashed unexpectedly. One of the benefits of NTFS was a transaction-tracking system, which made it possible for Windows NT to back out of any disk operations that were in progress when it crashed or lost power.

With NTFS, files, folders, and volumes can each have their own security. NTFS's security is flexible and built in. Not only does NTFS track security in ACLs, which can hold permissions for local users and groups, but each entry in the ACL can specify which type of access is given—such as Read & Execute, List Folder Contents, or Full Control. This allows a great deal of flexibility in setting up a network. In addition, special file-encryption programs were developed to encrypt data while it is stored on the hard disk.


Microsoft strongly recommends that all network shares be established using NTFS. Several current operating systems from Microsoft support both FAT32 and NTFS. It's possible to convert from FAT32 to NTFS without losing data, but you can't do the operation in reverse. (You would need to reformat the drive and install the data again from a backup tape.)



If you're using FAT32 and want to change to NTFS, the convert utility will allow you to do so. For example, to change the E: drive to NTFS, the command is `convert E: /FS:NTFS`.

Share permissions apply only when a user is accessing a file or folder through the network, as shown in Figure 6.7. *NTFS permissions* and attributes are used to protect the file when the user is local. With FAT and FAT32, you do not have the ability to assign “extended” or “extensible” permissions, and the user sitting at the console effectively is the owner of all resources on the system. As such, they can add, change, and delete any data or file.

FIGURE 6.7 Network share permissions and NTFS permissions



With NTFS as the filesystem, you are allowed to assign more comprehensive security to your computer system, as shown in Table 6.1. NTFS permissions can protect you at the file level. Share permissions can be applied to the folder level only, as shown in Table 6.2. NTFS permissions can affect users accessing files and folders across a network or logged in locally to the system where the NTFS permissions are applied. Share permissions are in effect only when the user connects to the resource through the network.

TABLE 6.1 NTFS permissions

NTFS permission	Meaning	Object used on
Full Control	Gives the user all the other choices and the ability to change permissions. The user can also take ownership of the folder or any of its contents.	Folder and file objects
Modify	Combines the Read & Execute permission with the Write permission and further allows the user to delete everything, including the folder.	Folder and file objects
Read & Execute	Combines the Read permission with the List Folder Contents permission and adds the ability to run executables.	Folder and file objects
List Folder Contents	The List Folder Contents permission (known simply as List in previous versions) allows the user to view the contents of a folder and to navigate to its subdirectories. It does not grant the user access to the files in these directories unless that is specified in file permissions.	Folder objects
Read	Allows the user to navigate the entire folder structure, view the contents of the folder, view the contents of any files in the folder, and see ownership and attributes.	Folder and file objects
Write	Allows the user to create new entities within a folder.	Folder and file objects

TABLE 6.2 Share permissions


Share permission	Meaning
Full Control	Gives the user all the other permissions as well as permission to take ownership and change permissions.
Change	Allows the user to overwrite, delete, and read files and folders.
Read	Allows the user to view the contents of the file and to see ownership and attributes.

Allow vs. Deny

Within NTFS, permissions for objects fall into one of three categories: Allow, Deny, or not configured. When viewing the permissions for a file or folder, you can select the box for

Allow, which effectively allows the group selected to perform that action. You can also deselect the box for Allow, which does not allow that group that action, as shown in Figure 6.8. Alternatively, you can select the Deny box, which prevents that group from using that action. There is a difference between not allowing (a cleared check box) and Deny (which specifically prohibits), and you tend not to see Deny used often. Deny, when used, trumps other permissions.

FIGURE 6.8 NTFS folder permissions




Permissions set on a folder are inherited down through subfolders unless otherwise changed. Permissions are also cumulative; if a user is a member of a group that has Read permission and a member of a group that has Write permission, they effectively have both Read and Write permissions.

Effective Permissions

When a user accesses a file share, both the share permissions and NTFS permissions interact with each other to form the effective permission for the user. Figure 6.9 shows that a user named Fred has logged in and received his access token containing the Sales and R&D groups, since he is a member of both groups. When Fred accesses the Sales file share, the share permissions define that he has read-only access because he is part of the Sales group. You can see that the NTFS permissions are granting him read and write access because of his

Sales group membership, as well as full control because he is also in the R&D group. If Fred were to locally log in to this computer, he would effectively have full control of these files. However, because he is accessing these files from the network, he only has read-only access because of the file-share permissions. The opposite is also true: if he had full permission at the share level and read-only permission at the NTFS level, he would effectively have read-only access.

FIGURE 6.9 Effective permissions



The rule for figuring out effective permissions is simple: If a user is in more than one group for which there are multiple permissions, take the most permissive permission of NTFS and then the most permissive permission of the share; the effective permission is the more restrictive of the two. There are some circumstances that change this rule slightly when the user (or group) is denied. If a user is in any group that is denied permission at the share or the NTFS level, they are denied for that access level. Therefore, when you derive the more restrictive permission, it will always be a deny for the user. A simple way to remember this is that a deny is a deny.

Moving vs. Copying Folders and Files

When you copy a file, you create a new entity. When you move a file, you simply relocate it and still have but one entity. This distinction is important when it comes to understanding permissions. A copy of a file will have the permissions assigned to it that are already in place at the new location of the file, regardless of which permissions were on the original file.

A moved file, on the other hand, will attempt to keep the same permissions as it had in the original location. Differences will occur if the same permissions cannot exist in the new location. For example, if you are moving a file from an NTFS volume to

FAT32, the NTFS permissions will be lost. If, on the other hand, you are moving from a FAT32 volume to an NTFS volume, new permissions will be added that match those for newly created entities.



Folder copy and move operations follow guidelines that are similar to those for files.

File and Folder Attributes

Since the introduction of FAT, file and folder *attributes* have existed in the filesystem. The basic set of attributes—Read-only, Hidden, System, and Archive—can still be found in both FAT and NTFS and are very useful for the operating system, as explained in the following list. You can view the file or folder attributes by right-clicking the file or folder and selecting Properties. From there, you can view the Read-only and Hidden attributes.

Read-Only The Read-only attribute is used to make a file or folder read-only. If the Read-only attribute is present, regardless of the permission of the file or folder, the user cannot write or delete a file or folder. The Read-only attribute is useful for protecting files from being overwritten or deleted.

Hidden The Hidden attribute makes files and folders disappear when using the Windows File Explorer utility, using the defaults or using the `dir` command. The files still function and are still a part of the filesystem, but they just don't appear by default. The Hidden attribute is useful when you want to hide a file or folder from the average user.


System The System attribute marks files and folders as system files. When a file or folder is marked as system, it means the file or folder is critical. With the introduction of NTFS, the System attribute is not really useful anymore, but it is still there for backward compatibility.

Archive The Archive attribute can be used to tell the system whether the file has changed since the last time it was backed up. Technically, it is known as the Archive Needed attribute. If this bit is on, the file or folder should be backed up. If it is not selected, a current version of the file is already backed up.



Only the Read-only and Hidden attributes can be set or cleared from the GUI. The Archive attribute can be accessed from the Advanced Attributes dialog box. The System attribute can be set or cleared only with the `attrib.exe` command.

Windows uses NTFS, which gives you a number of options that are not available on earlier filesystems, such as FAT and FAT32. A number of these options are implemented through the use of the Advanced Attributes dialog box, as shown in Figure 6.10.

FIGURE 6.10 The Advanced Attributes dialog box

FAT32 does not have as many options as NTFS, such as encryption and compression. These attributes are available only on NTFS partitions.

To reach these options in Windows, right-click the folder or file that you want to modify, and then select Properties from the menu. On the main Properties page of the folder or file, click the Advanced button in the lower-right corner. In the Advanced Attributes window, you have access to the following settings:

Indexing Windows implements a feature called the *Indexing Service* to catalog and improve the search capabilities of your drive. Once files are indexed, you can search for them more quickly by name, date, or other attributes. Setting the index option on a folder causes a prompt to appear, asking whether you want the existing files in the folder to be indexed as well. If you choose to do this, Windows will automatically reset this attribute on subfolders and files. If not, only new files created in the folder will be indexed.

Compression The versions of Windows that you need to know for the exam support advanced *compression* options, which were first introduced in Windows NT. NTFS files and folders can be dynamically compressed and uncompressed, often saving a great deal of space on the drive. As with indexing, when you turn on compression for a folder, you'll be prompted as to whether you want the existing files in the folder to be

compressed. If you choose to do this, Windows automatically compresses the subfolders and files. If not, only new files created in the folder are compressed.



Compression works best on such files as word processing documents and uncompressed images. Microsoft Word files and Microsoft Paint bitmaps can be compressed up to 80 percent. Files that are already packed well do not compress as effectively; EXE and ZIP files generally compress only about 2 percent. Similarly, GIF and JPEG images are already compressed (which is why they are used in Internet web pages), so they compress a little or not at all.

Encryption *Encryption* lets you secure files so that no one else can view them. You encrypt files by encoding them with a key to which only you have access. This can be useful if you're worried about extremely sensitive information, but in general encryption is not necessary on the network. NTFS local file security is usually enough to provide users with access to what they need and prevent others from getting to what they shouldn't. If you want to encrypt a file, go through the same process that you would for indexing or compression.



Encryption and compression are mutually exclusive—you can set one but not both features on a file or folder. Not all features are available in all editions of every operating system. If a user forgets their password or is unable to access the network to authenticate their account, they will not be able to open encrypted files. By default, if the user's account is lost or deleted, the only other user who can decrypt the file is the Administrator account.

In Exercise 6.2 you will have the opportunity to view file permissions for both basic and advanced permissions. Make sure that you don't inadvertently add any deny permissions, as you could be prevented from making any further changes.

EXERCISE 6.2

Examining File Permissions

1. Open Windows File Explorer.
2. Right-click a file or folder and choose Properties.
3. Select and then examine the Security tab.

You'll see the users and/or groups to which permissions have been assigned.

4. Select a user or group in the list, and examine the list of standard permissions. (To add a new user or group, click Add and follow the prompts.)

EXERCISE 6.2 (continued)

Any standard permissions that are checked in the Allow column are applied. If a check box is grayed out, then the permission was inherited.

5. To revoke a set of standard permissions, click the appropriate check box in the Deny column.

If you click the check box in the Deny column for the Full Control permission, all other standard permissions are also denied.

6. Click Advanced to examine advanced options.
7. Click Cancel twice to close the file or folder's properties.

Be sure that you don't accidentally make any changes that you didn't intend to make. Changing permissions without understanding the ramifications can have negative consequences, such as losing access to files or folders. It is a best practice to assign Deny permissions sparingly. It's better to uncheck Allow. (You may need to turn off Inheritance.)

Shared Files and Folders

You can share folders, and the files beneath them, by right-clicking the file or folder and selecting Show More Options, then choosing Give Access To from the context menu and selecting Specific People. In Windows, the context menu asks you to choose with whom you want to share the folder or file, as shown in Figure 6.11. You can then choose who to share it with along with their respective permissions. It is important to understand that when you use this method to share files and folders, the share permissions are set to Full Control for the Everyone group. The dialog box shown in Figure 6.11 will allow you to manipulate the NTFS permissions via the permission level.


You can access the Advanced Sharing settings by right-clicking the folder you want to share, selecting Show More Options, then Properties, then clicking the Sharing tab, and finally selecting Advanced Sharing, as shown in Figure 6.12. This file-sharing method is more traditional with network administrators because every aspect of the share can be controlled. Using this method, only the share permissions are set from this dialog box. The NTFS security permissions are set on the Security tab. In addition, you can add other share names to the same location, limit the number of simultaneous connections, and add comments.

Administrative Shares vs. Local Shares

Administrative shares are automatically created on all Windows operating systems on the network for administrative purposes. These shares can differ slightly based on which operating system is running, but they always end with a dollar sign (\$) to make them hidden. There is one share for each volume on a hard drive (c\$, d\$, and so on) as well as admin\$

(the root folder—usually C:\WINDOWS) and print\$ (where the print drivers are located). These shares are created for use by administrators and usually require administrator privileges to access.

FIGURE 6.11 Choose People To Share With



Local shares, as the name implies, are shares that are created locally by the administrative user on the operating system. The term *local shares* is used to distinguish between automated administrative shares and manually created shares.

Permission Inheritance and Permissions Propagation

Inheritance is the default throughout the permission structure unless a specific setting is created to override it. A user who has Read and Write permissions in one folder will have them in all the subfolders unless a change has been made specifically to one of the subfolders. If a user has the Write permission, which is inherited from the folder above, the removal of permissions cannot be performed unless inheritance is disabled. Only additional permissions can be added explicitly, since this is actually a new permissions entry and not the removal of an existing permissions entry. You can control NTFS inheritance by right-clicking a folder, selecting Show More Options, then selecting Properties, then choosing the Security tab, selecting Advanced, and clicking the Disable Inheritance button, as shown in Figure 6.13.

FIGURE 6.12 Advanced file and folder sharing





FIGURE 6.13 Disabling inheritance




If the Disable Inheritance button is selected, you will be changing the NTFS security settings on the folder. A second dialog box will pop up and you must choose to either keep the existing permissions, by selecting Convert Inherited Permissions To Explicit Permissions On This Object, or to start fresh with no permissions by selecting Remove All Inherited Permissions From This Object. You should use caution when selecting the latter of the two options, because all folders and files below will inherit and propagate the removal of existing permissions.

If you want to make sure that inheritance and permissions for a folder are propagated to all files and folders below, you can use Replace All Child Object Permission Entries With Inheritable Permission Entries From This Object (refer to Figure 6.13). This option will replace every permission in this folder and all the subfolders, regardless of whether explicit permissions were applied further down in the folder structure.

In the Advanced Security Settings, you can also configure permissions entries that only apply to the current folder, current folder and files, all folders and files, or other variations of these, as shown in Figure 6.14. These settings can change the propagation of file permissions to folders and files.

FIGURE 6.14 Permission entry



System Files and Folders

System files are usually flagged with the Hidden attribute, meaning they don't appear when a user displays a folder listing. You should not change this attribute on a system file unless absolutely necessary. System files are required in order for the operating system to function. If they are visible, users might delete them (perhaps thinking that they can clear some disk

space by deleting files that they don't recognize). Needless to say, that would be a bad thing! Most system files and folders are protected by the operating system and won't allow deletion, but better safe than sorry.

Changing File Attributes


File attributes determine what specific users can do to files or folders. For example, if a file or folder is flagged with the Read-only attribute, then users can read the file or folder but cannot make changes to it or delete it. Attributes include Read-only, Hidden, System, and Archive, as well as Compression, Indexing, and Encryption. Not all attributes are available with all versions of Windows. We'll look at this subject in more detail in a moment.



Some attributes—such as Read-only, Hidden, System, and Archive—date back to DOS. All others—such as Compression, Indexing, and Encryption—are a part of NTFS.

You can view and change file attributes either by entering **attrib** at the command prompt or by changing the properties of a file or folder. To access the properties of a file or folder, right-click the file or folder and select Show More Options, then Properties. You can view and configure the Read-only and Hidden file attributes on the General tab. To view and configure additional attributes, click Advanced, as shown in Figure 6.15.

FIGURE 6.15 Windows file attributes



Windows Security Features

Although everything we've covered so far is a Windows security feature, we have focused on the basic elements of authentication and securing resources. In the following section we will cover the common Windows security features that can protect us from external threats, such as malware, malicious network activity, and data loss.

Microsoft Defender Antivirus

Microsoft Defender, also known as Windows Defender Antivirus, was originally introduced with Windows XP as a downloadable antivirus product. It was later shipped with the Vista operating system, and it has become a pillar of security in the Windows 11 operating system. The addition of Microsoft Defender allows the Windows operating system to be protected right out of the box. A user never needs to install anything else to be protected from malware. That being said, there are lots of products on the market that provide security features above and beyond Microsoft Defender. You can view the Microsoft Defender settings by navigating to Start > Settings > Privacy & Security > Windows Security > Virus & Threat Protection in Windows 11, as shown in Figure 6.16.



You may have noticed that the Microsoft Defender Antivirus branding is slowly disappearing from Windows 11. The Microsoft Security app is slowly replacing the branding of Microsoft Defender Antivirus. There are a few reasons for this; the first is Microsoft Defender Antivirus is being sold as an add-on product for businesses. Another reason is the Microsoft Security app does more than just antivirus, malware control, ransomware mitigation, and a wealth of other security-related features discussed throughout this book. The Microsoft Defender Antivirus product is still under the hood protecting the operating system right out of the box. This book will reference the Microsoft Defender Antivirus product since the CompTIA objectives still reference the product. More information on the Microsoft Security app and Microsoft Defender Antivirus can be found at <https://learn.microsoft.com/en-us/windows/security/operating-system-security/system-security/windows-defender-security-center/windows-defender-security-center>.

On the Virus & Threat Protection screen you can initiate a scan of the computer. In addition, the screen will detail how many threats have been found, when it was last scanned, how many files were scanned, and how long the scan lasted. By clicking Scan Options you can select from Quick Scan, Full Scan, Custom Scan, and Microsoft Defender Offline Scan.


Quick Scan This option will check folders in the system where threats are commonly found.

Full Scan This option will perform a full scan of the entire operating system.

Custom Scan This option will allow you to select the folders on the drive you want to scan.

Microsoft Defender Antivirus (Offline Scan) This option will restart the device, reboot into the Windows Recovery Environment (WinRE), and launch a full Microsoft Defender Scan. It is used to remove malware that is difficult to remove while the operating system is running.

FIGURE 6.16 Microsoft Defender settings




You can also change the Microsoft Defender Virus & Threat Protection Settings by clicking Manage Settings on the Virus & Threat Protection screen. This will allow you to manage a number of settings to change the way Microsoft Defender operates, as shown in Figure 6.17.

Real-Time Protection You can toggle off real-time protection when installing certain applications that require that antivirus be off during installation. However, the real-time protection will turn back on automatically after a period of time.

Dev Drive Protection Dev Drive Protection in Microsoft Defender is a feature designed to enhance performance for developers using Windows 11. When enabled, Dev Drive Protection runs in a special asynchronous performance mode, which defers security scans until after file operations are completed. This mode improves system performance without compromising security.

FIGURE 6.17 This screen provides you with several options to specify how protection is handled by Microsoft Defender Virus & Threat Protection.



Cloud-Delivered Protection You can also toggle Cloud-Delivered Protection, which provides cloud-based data on threats and ultimately faster protection. Turning this setting off might be required for certain regulatory requirements, since it automatically turns on cloud-based sample submission.

Automatic Sample Submission Automatic Sample Submission allows Microsoft Defender to automatically submit samples to Microsoft for analysis. Automatic Sample Submission can be toggled on and off.

Tamper Protection The Tamper Protection security setting prevents malicious applications from tampering with Microsoft Defender settings. Tamper Protection protects against tampering from third-party processes; even Group Policy settings cannot disable Microsoft Defender when Tamper Protection is turned on.

Controlled Folder Access The Controlled Folder Access feature can be accessed in the Virus & Threat Protection Settings. The feature is used to prevent ransomware from harming files, folders, and memory areas on the device. By default, it is turned off, but

it can be easily turned on by clicking the toggle. Once it is set to the default, folders protected are documents, pictures, videos, music, and favorites. Specific folders can be added, and the defaults can also be removed. You can also exempt an application so that it is allowed to modify the files.

Exclusions Although it is not common to need to exempt a folder from the antimalware engine, it can be configured from the Virus & Threat Protection Settings. This option allows you to exempt an entire folder from the real-time protection and scans. One use may be performance-impacted applications such as games, but caution should be used.

Update Definitions An antimalware engine is only as good as its latest definitions, and Microsoft Defender is no different. Therefore, both the Microsoft Defender engine and its definitions are updated quite frequently. Both are updated through the Windows Update process.

Windows Defender Firewall

Windows Defender Firewall is an advanced host-based firewall that was first introduced with Windows XP Service Pack 2. It was integrated and became a security feature with the introduction of Windows Vista. While host-based firewalls are not as secure as other types of firewalls, Windows Defender Firewall provides much better protection than in previous versions of Windows, and it is turned on by default. Windows Defender Firewall is used to block access from the network, which significantly reduces the surface area of attack for the Windows operating system.


To access Windows Defender Firewall in Windows 11, navigate to Start > Settings > Privacy & Security > Windows Security > Firewall & Network Protection. Windows Defender Firewall is divided into separate profile settings: for domain networks (if you're connected to a domain), private networks, and public networks. In Figure 6.18, you can see the default protection for a Windows client that is not joined to a domain and is active on a public network.

Activate and Deactivate the Firewall By default, the Windows Defender blocks inbound connections, but outbound connections are not blocked by default. You can temporarily deactivate the Windows Defender Firewall for a network profile by clicking on the network profile in the Firewall & Network Protection screen and toggling the Windows Defender Firewall setting to Off, as shown in Figure 6.19. Deactivating the firewall protection is extremely useful if you are trying to troubleshoot a network connectivity problem. Dropping the firewall temporarily and testing the incoming connection will confirm that a firewall adjustment needs to be made.

Allowing Applications When applications attempt to listen on a port for an incoming connection, the Windows operating system will display a notification asking if you want to allow access. However, if the notification does not display and you determine that an application needs to be allowed through the firewall, you can click Allow An

App Through Firewall in the Firewall & Network Protection screen. This will open the Allowed Apps settings page shown in Figure 6.20. Here you can modify the apps that are allowed through the Windows Defender Firewall, as well as add new ones. You will need administrative access to change any network firewall settings.

FIGURE 6.18 Firewall & Network Protection



Windows Defender Firewall with Advanced Security When you allow an application to listen for an incoming connection via the notification dialog box, the operating system creates a rule in the firewall to allow the connection. This is all done for you behind the scenes and it shows up neatly as an allowed application. However, you can also manually create a rule in Windows Defender Firewall with the Advanced Security MMC, as shown in Figure 6.21. You can open the MMC by clicking Advanced Settings on the Firewall & Network Protection screen.

FIGURE 6.19 Deactivating Windows Defender Firewall network protection




FIGURE 6.20 Windows Defender Firewall Allowed Apps






FIGURE 6.21 Windows Firewall with Advanced Security

Here you can configure inbound and outbound rules as well as import and export policies and monitor the security of your system. Monitoring is not confined to the firewall; you can also monitor security associations and connection security rules. In short, Windows Defender Firewall with Advanced Security is an incredibly powerful tool that builds on what Windows Vista introduced. Not only can this MMC snap-in do simple configurations, but it can also configure remote computers and work with Group Policy.

Exceptions Manual exceptions, also known as firewall rules, are configured on the Inbound Rules tab in the Windows Defender Firewall with Advanced Security MMC. You can click **New Rule** and configure a firewall rule based on a program, port, predefined rule, or totally custom rule. A rule based on a program allows all incoming connections to the application. A rule based on a port allows you to configure a specific rule based on a TCP or UDP connection to a specific port or range of ports. A predefined rule allows you to modify a predefined rule. A custom rule allows you to configure a program and specific ports; you can even scope it down to the incoming range of IP addresses you will allow. Any of these settings can also be configured after running the New Rule Wizard, as shown in Figure 6.22.

FIGURE 6.22 Windows Defender Firewall with Advanced Security inbound rules

BitLocker

You have to be careful, because CompTIA sometimes refers to the utility as “bit-locker” or “Bitlocker,” while it is officially known as *BitLocker*. This tool allows you to use drive encryption to protect files—including those needed for startup and login. This is available only with more complete editions of Windows 11 (Pro, Enterprise, Education, Pro for Workstations).


Another requirement is the use of a *Trusted Platform Module (TPM)*. The TPM is a chip on the motherboard that safely stores the encryption key so that the key is not stored on the encrypted disk. BitLocker can mitigate the risk of data loss, because if the disk is separated from the computer, it is still encrypted and only the TPM with the encryption/decrypting key can decrypt the disk. This prevents out-of-band attacks on the hard drive, where it would be mounted and examined on a second system. BitLocker can also sense tampering. If it senses tampering, the recovery key must be reentered. The recovery key is either entered from a printout, loaded from a USB drive in which it was originally saved, or recovered from your Microsoft account. An option of how the recovery key is stored is presented to you when you initially turn on BitLocker.

BitLocker to Go

You can also protect removable drives with *BitLocker to Go*. It provides the same encryption technology BitLocker uses to help prevent unauthorized access to the files stored on them. You can turn on BitLocker to Go by inserting a USB drive into the computer and

opening the BitLocker Drive Encryption Control Panel applet, as shown in Figure 6.23. When a USB drive is inserted into a Windows computer that contains BitLocker to Go encryption, the operating system prompts you for the password to unlock the drive. This password is the one you used originally when you set up BitLocker to Go on the USB drive.


FIGURE 6.23 BitLocker Drive Encryption applet



EFS

Encrypting File System (EFS), available in most editions of Windows, allows for the encryption/decryption of files stored in NTFS volumes. EFS uses certificates to encrypt the data, and the private certificate is stored in the user profile. When the first file is encrypted, the operating system automatically generates a key pair. If the computer were joined to an Active Directory domain and a *certificate authority (CA)* existed, the CA would create the key pair. You can encrypt a file or folder by right-clicking the object, selecting Show More Options, then Properties, then Advanced, as shown in Figure 6.24.

All users can use EFS, whereas only administrators can turn on BitLocker. EFS does not require any special hardware, whereas BitLocker benefits from having the TPM. As an additional distinction, EFS can encrypt just one file, if so desired, whereas BitLocker encrypts the whole volume and whatever is stored on it. Finally, EFS can be used in conjunction with BitLocker to further increase security.

FIGURE 6.24 Encrypting a file in Windows 11

Exam Objective 2.11, “Given a scenario, configure relevant security settings in a browser,” requires you to know how to configure security settings in a browser. In addition, you should know about browser installation, patching, extensions and plug-ins, password managers, the various settings, and browser feature management.

Web Browser Security


The web browser is arguably the most used application on the Windows operating system. It is also your portal to the Internet and all the bad things that reside outside of your network. Therefore, it makes sense that there is an entire objective dedicated to web browser security for the CompTIA exam. In this section we will explore various browser security topics. We will primarily focus on the Microsoft Edge browser that comes preinstalled in the Windows 11 operating system. However, we will also reference Chrome, as it holds more than 68 percent of the market share worldwide (<https://gs.statcounter.com>) as of this writing.

Browser Download and Installation

Edge is the successor to Internet Explorer and comes preinstalled in the Microsoft Windows 11 operating system. Edge is also the default web browser in the Windows 11 operating system. Although Edge comes preinstalled in Windows 11, there are circumstances when you need to download the Edge web browser and install it in other operating systems. A common scenario is downloading and installing Edge for a server operating system in which it is not preinstalled, such as Windows Server 2019 or another operating system, such as macOS or Linux.

To download the Edge installer, you can often just search for “Edge download.” Depending on the search you use, the result can be in the top three results, but malware can also be disguised as the download. Therefore, you need to ensure your download is supplied by a trusted source. You should only download the browser from the vendor, in this case Microsoft. A simple way to check that the source is the vendor is to check the URL in the address bar, as shown in Figure 6.25. If the address bar shows the parent address of `https://www.microsoft.com`, then you are obviously downloading Edge from a trusted source.

FIGURE 6.25 Downloading Microsoft Edge



There are two main ways to download most browsers: online and offline. The online version will initially download a small install application (approximately 2 MB) that will download the rest of the installation and install the remaining download. The online version typically can be downloaded from the Microsoft Edge download page, as previously described. An offline installation can be downloaded from Microsoft as well and offers the benefit on not requiring any Internet connectivity. An offline version can be downloaded from www.microsoft.com/en-us/edge/business/download. By downloading the offline version from the Microsoft trusted source, you can be assured that future installations from this install are genuine.

When you store the installation for future installation, you should protect the installer from being tampered with. The best preventive measure is to create a Secure Hash Algorithm (SHA) hash of the executable and store it in another location. Before you run the offline installer, simply run the hashing against the install and check its signature. If the signature matches, then the installer has not been tampered with. If the signature does not match, tampering could have occurred, and it should be treated as being from an untrusted source.

In the following exercise, you will get to use the `Get-FileHash` cmdlet built into PowerShell. If you already have a file hash and it is not in the SHA256 format, you can use the argument of `-Algorithm MD5` or `-Algorithm SHA1` depending on the format you need to verify.

In Exercise 6.3 you will create a hash for a sample file with a PowerShell cmdlet. Then you will slightly modify the file and hash it again and compare the outcomes.

EXERCISE 6.3

Working with File Hashes

1. Open Notepad, type **example**, and save it to the desktop as **example.txt**.
 2. Navigate to the desktop, hold the Shift key, right-click, and choose Open PowerShell Window Here.
 3. Type **Get-FileHash .\example.txt**.
 4. Examine the file hash and make a note of it.
 5. Type **Get-FileHash .\example.txt**.
 6. Compare the two file hashes; they should be identical.
 7. Open the file `example.txt` and make the lowercase `e` in `example` an uppercase `E`; then save the file.
 8. Navigate back to the PowerShell prompt.
 9. Type **Get-FileHash .\example.txt**.
 10. Examine the file hash.
 11. Change the uppercase `E` back to a lowercase `e`, rerun the `Get-FileHash` cmdlet, and compare the two hashes.
-

Google Chrome and other third-party browsers require download and installation from the vendor sites, just like Microsoft Edge. These practices can and should be employed with any installation, not just web browsers. You should always download installations from a trusted source, then protect them with file hashing to detect tampering.

Browser Patching

A serious security concern is the number of web browser exploits released every day. A malicious web page could automatically download malware or allow cross-site scripting and compromise your login. These are just a few exploits that commonly need to be patched on a consistent basis. Therefore, constant browser patching is necessary. The version you download today is not going to be the version you will have tomorrow.

Each web browser handles browser patching similarly. Chrome employs a silent update mechanism, which automatically installs updates without user intervention. This method ensures that the browser remains secure and up-to-date. This approach minimizes the risk of vulnerabilities being exploited by attackers. Microsoft Edge, built on the Chromium engine, also utilizes an auto-update feature but offers additional flexibility for enterprise environments. An administrator can use patching tools like Microsoft Configuration Management to deploy constant updates.

When a web browser update is available, the user typically receives a notification indicating that a new version is ready to be installed. Inside Chrome this notification might appear as a subtle change in the browser's menu icon, prompting users to relaunch the browser to apply the update. Users can choose to ignore this notification temporarily, but they will continue to see reminders until they restart the browser. Expecting the user to relaunch the browser is fine in some cases. However, an administrator can enforce a policy that requires users to relaunch the browser within a specified time frame. Once the user decides to proceed with relaunching, the browser will close and reopen, applying the update.


If a system reboot is necessary, the user will be prompted to restart their device to complete the update process. This is typically required when Windows Update performs major upgrades of Edge periodically.

Extensions and Plug-Ins

The best feature in any web browser is the ability to extend its functionality to accommodate your exact needs. By using extensions and plug-ins, you can functionally change the way the web browser works and the web pages it consequently renders. *Extensions* extend the functionality of the web browser that was not originally conceived when it was designed. *Plug-ins* change the way web pages are rendered on the web browser. Each web browser will call these something slightly different—for example, plug-ins might be called add-ons, as in the case of Firefox. Google Chrome refers to both extensions and plug-ins as extensions. Microsoft refers to both extensions and plug-ins as add-ons. For the remainder of this section, we will refer to all plug-ins and extensions as add-ons.

Most modern web browsers have an ecosystem of add-ons. Google Chrome has an ecosystem called the Chrome Web Store that can be accessed at <https://chrome.google.com/webstore>. Microsoft has an ecosystem called Edge Add-ons that can be accessed at <https://microsoftedge.microsoft.com/addons/microsoft-edge-extensions-home>, as shown in Figure 6.26. An add-on ecosystem is a place where the vendor trusts the publisher of an add-on. The add-on ecosystem in turn distributes the add-on for users, and supplies updates and synchronized installations across multiple devices. The add-on ecosystem should always be considered a trusted source for web browser add-ons.

FIGURE 6.26 Microsoft Edge add-ons




Most web browsers will also allow manual installation of add-ons. This is also called *sideloading*, since you are manually installing the file outside of the ecosystem. This is usually done and allowed by the web browser for development purposes. This type of installation can be found among untrusted applications that the ecosystem doesn't allow. The add-on might look reputable and the web page might explain that manual installation is required, since the developers don't agree to the terms of service (ToS) of the ecosystem. In any case, you should consider these add-ons untrusted and therefore avoid them.

Credentials Managers

As you sign up for websites, they require more complex, lengthier passwords. Sometimes your username might be your favorite nickname, sometimes the nickname might not be available, and sometimes the site might require an email in lieu of a nickname-style username. You should also employ the practice of single-purpose passwords, to prevent credential stuffing attacks, and never reuse a password. So, if your credentials on site A are compromised, then your credentials on site B will not be vulnerable.

No longer can we keep track of all these usernames and passwords. Luckily password managers have come to our rescue. They are built into every operating system and most web browsers. Microsoft Edge and the older Internet Explorer both use the Microsoft Credential Manager, which was originally introduced with Windows XP. It functions as a password manager and is built into the operating system. In Figure 6.27, you can see some web credentials stored by Microsoft Edge. You can access the Credential Manager by navigating to the Start menu, typing **Control Panel** and selecting it in the results, then clicking Credential Manager.

FIGURE 6.27 Microsoft Credential Manager



Credentials are stored by successfully logging into a website with a username and password combination. The web browser will ask if you want to save the credentials. Once the credentials are stored, when a website asks for a username and password matching the site in the Credential Manager, the associated credentials are offered to the user for logging into the site. If you are in the Credential Manager and you want to see the password, click Show and enter your credentials for the currently logged-on user. By entering your credentials for the currently logged-on user, you unlock the Credential Manager and it will let you see the password. You also have the option of deleting the credentials by clicking Remove.

Other web browsers, such as Google Chrome and Mozilla Firefox, have a credential manager built-in. These web browsers do not use the Microsoft Credential Manager. They synchronize their usernames and passwords between installations of the web browser on various devices. You can also download and install a third-party stand-alone credential manager. KeePass is one example, and there are several freely available on the Internet.


Secure Data Transfers

Securing data transfers to and from the web browser is critical for the security of day-to-day operations. We use our web browsers for accessing sensitive information at work, doing our personal banking, accessing social media, and the list goes on. So, securing data in transit is a primary concern to manage the risk of eavesdropping from threat agents. The Hypertext Transfer Protocol over Secure Sockets Layer (SSL), also known as HTTPS, is a method for securing web data communications. SSL is a cryptographic suite of protocols that use public key infrastructure (PKI) to provide secure data transfer.

A PKI is a system that provides the key pair of private and public keys, also known as certificates. The certificates are used to validate communication and encrypt communications between web browsers and web servers. The private key (certificate) from the key pair is installed on the web server. The public key (certificate) is available to anyone who wants to validate the data encrypted with the private key that is installed on the server. The web browser will automatically download the public key for the cryptography process.

Because there is a level of trust involved with the public/private key pair, your web browser must initially trust the publisher of the key pair. The publisher of the key pair is known as the certificate authority (CA). Every web browser comes with an initial list of trusted root certificate authorities, which is a beginning of trust for issuing CAs. The issuing CAs are the CAs that actually issue the key pairs. In most cases the web browser will use the operating system's list of trusted root CAs, as shown in Figure 6.28. However, some web browsers maintain a list of their own trusted root CAs.

Each certificate has an expiration date, and from time to time, certificates expire, get replaced, or in some cases get revoked for various reasons. Therefore, the list of trusted root certificate authorities must be updated now and then. The Microsoft Windows platform does this with routine Windows Updates, but browsers may update their own lists with updates of their own.

FIGURE 6.28 Trusted root CAs

Although the S in HTTPS stands for secure, SSL is the protocol suite used for securing communications. The SSL suite is a suite of protocols that includes the current standard of Transport Layer Security (TLS) 1.3. TLS 1.3 is used for securing websites, as of this writing.

Settings

A lesson on web browser security wouldn't be complete without covering security settings for the web browser. Next, we will explore the various security settings you will find on modern-day web browsers. We discuss the most common topics that you will see on the CompTIA exam, but in no way is this a complete list of security settings. The web browser is the most used application for accessing data, and new settings are introduced all the time to combat threat agents.


Pop-Up Blocker

Pop-up blockers are used to prevent both pop-ups and pop-unders from appearing when you visit a web page. While older browsers did not incorporate an option to block pop-ups, most current browsers, including the latest versions of Edge and Chrome, have that capability built in.

By default, the pop-up blocker is enabled on all sites globally. However, some sites, such as banking websites, might need to pop up a web page. When this happens, the

address bar will notify you that an attempted pop-up occurred. You can then click and allow the specific site to use pop-ups. If you want to view a site's permission for pop-ups, click the lock at the left of the address bar and choose Permissions For This Site, as shown in Figure 6.29.

FIGURE 6.29 Permissions for a site




Clearing Browsing Data

Browsing data is a broad term used to describe any data stored while visiting websites. Data that may be considered browsing data is browsing history, download history, cookies, image and file caches, passwords, autofill form data, and site settings. These are just the top categories, so keep in mind that browsing data is a broad term. Browsing history is used by the browser to highlight links you've already visited. Download history allows you to historically see what has been downloaded. Cookies are used to save settings used by the web page, such as themes and login data, to name a few. Passwords are saved for convenience when logging into websites, the same as autofill form data. The site settings are adjustments you've made to the browser to optimally render the web page.

There are various reasons you may want to clear browsing data. The most compelling reason is data privacy. However, it is sometimes necessary to clear browsing data when you are trying to replicate a problem seen with the web browser.

You can clear the browsing data by clicking the three dots in the upper-right corner of the Edge web browser. Then click Settings, choose Privacy, Search, And Services, and scroll down to Delete Browsing Data Now and select Choose What To Clear. You will be presented with a dialog box similar to Figure 6.30. Options for the time range are Last Hour, Last 24 Hours, Last 7 Days, Last 4 Weeks, and All Time. You can also selectively delete the web browsing data that you desire.

FIGURE 6.30 Clearing browsing data



Clearing Cache

When a web browser renders a web page, the files retrieved are cached. This is done so that if you need them again you can quickly retrieve them from storage. This caching mechanism speeds up the web browser and reduces unneeded trips to the Internet. There are times when you need to clear your web browser cache, such as when developing a web page. You will


want to retrieve the latest copy of the web page and its assets so that you can verify how it is rendered. The cache images and files are part of the web browsing data that can be cleared. The process is similar to the previously mentioned process for clearing browsing data, except only the cache images and files will be cleared.

Private-Browsing Mode

Private-browsing mode was created primarily to address data privacy while you are web browsing. The most compelling feature of private-browsing mode is that it does not store any web browsing data. Therefore, when you close private-browsing mode, all browsing data is destroyed.

The private-browsing mode on the Edge web browser is called InPrivate browsing; Chrome calls it Incognito mode. On the Edge web browser you can enter InPrivate mode by clicking the three dots in the upper-right corner of the window, then selecting New InPrivate Window. Another way is to press Ctrl+Shift+N to open the InPrivate window. In either case, the window explains the mode, as shown in Figure 6.31. The methods to open a private-browsing window on Chrome are identical, except that it is called Incognito mode.


FIGURE 6.31 InPrivate browsing



Browser Data Synchronization

On average, we access our web browsers on several devices. You may access them on your home computer, personal laptop, and mobile device, just to name a few. Therefore, the browsing data should follow you, no matter which device you are using. Modern web browsers fortunately support browser data synchronization, which allows you to share your data across many different devices. To sync your device, you just need to make sure that you have signed in and acknowledged that you want to sync browsing data. You can verify this by clicking on your account in the upper-left corner of the Edge web browser. Your account details will be displayed, as shown in Figure 6.32. If you are not syncing, you can simply sign in.

FIGURE 6.32 Account details



Use caution before syncing your personal web browsing data with your work computer's. You should keep a strict separation between work and leisure. During our leisure time we might search for something that could be misconstrued as mischievous to our coworkers. Therefore, you should maintain different accounts for work and play and never sync your personal account to a work web browser.



Real World Scenario

Leisure vs. Work Browsing

During leisure time you may be watching a movie in which someone picks a lock and curiosity may strike you to search “how to pick a lock.” That search will now be in your autofill form data (browsing data). Imagine clicking a search box while giving a presentation in front of coworkers and “how to pick a lock” autofills. You know the context of your search, but do your audience and coworkers know? Think of all the other searches you perform in your leisure time.

Ad Blockers

The Internet was originally embraced as a mechanism to share information and bring people together. However, it was equally embraced by commerce as a mechanism to deliver products and services. One such service is the marketing of other services and products; we get these as marketing ads.

You can install an ad blocker add-on for most web browsers to stop a lot of the spammy ads you might receive on a web page. However, entire websites, such as many news sites, subsidize their income with marketing ads. If you don't have a subscription, then you have to allow marketing ads in order to enjoy their content; it seems a pretty fair trade. Most of these sites allow you to have directions on how to exempt their site from the ad blocker.

You may wonder why you should install an ad blocker at all if every site allows you to exempt it. Ad blockers are extremely useful as an added layer of insulation from threat actors. Although most sites are legitimate and serve relevant ads, there are plenty of sites that serve malicious ads. An ad blocker helps to block all ads and allows you to judge whether the site is worthy of an exemption.


Proxy

A web browser can be configured to use a proxy server, which fetches the request on behalf of your request and returns the data. Proxy servers are also effective for controlling policies of what an employee can view. This provides not only a layer of protection, but also a way to centrally log requests. Filtered access to certain web sites is something that is decided at

the organizational level. It is common to block gambling, pornography, and other illegal or illicit sites.

The proxy setting can be configured for most web browsers at the operating system level. This can be accomplished by opening the Start > Settings > Network & Internet > Proxy, as shown in Figure 6.33. From the proxy screen you can manually configure a proxy server, or allow the automatic detection through the use of JavaScript. It is common for an organization to policy the proxy server settings, so they are automatically applied.

FIGURE 6.33 Proxy Settings



Secure DNS

Secure DNS, also known as DNS over HTTPS (DoH), is a technology designed to enhance privacy and security on the Internet. The need for Secure DNS arises from the vulnerabilities associated with traditional DNS queries, which are typically sent in plain text. This makes them susceptible to eavesdropping, manipulation, and other cyberthreats. By encrypting DNS queries, Secure DNS ensures that these lookups are protected from interception and tampering by malicious actors. This in turn safeguards the user's browsing activities and personal information.

Configuring Secure DNS in Chrome is performed by navigating to the browser's settings, selecting Privacy And Security, and then Security. Here, they can find the Use Secure DNS option, which allows them to choose a preferred DNS provider or enter a custom provider. Similarly, Microsoft Edge offers an easy way to enable Secure DNS. Users can go to Settings, then Privacy, Search, and Services, and under the Security section, they can turn on Use Secure DNS To Specify How To Look Up The Network Address For Websites. Edge also provides options to use the current service provider or select a custom provider. Secure DNS is enabled in both Edge and Chrome by default.

Browser Features

Web browser features can be significantly extended through the installation of extensions, add-ons and plug-ins, which are small software modules designed to add functionality and enhance the user's browsing experience. Extensions, also known as add-ons, modify the core functionality of the browser, allowing users to customize their browsing environment with tools like ad blockers, password managers, and note-taking utilities. Plug-ins, on the other hand, provide additional capabilities without altering the browser's core functions, such as enabling PDF editing or media playback. Their differentiations are so blurred that most web browsers call them add-ons or extensions; these words are interchangeable.

These add-ons are typically available through browser-specific web stores, where users can browse, read reviews, and install them with a simple click. By leveraging these add-ons, users can tailor their browsers to meet their specific needs, improving productivity, security, and overall user experience.



Management

When you want to add an add-on to your web browser, you will typically visit the web browser's add-ons website. This is how you will get new features for your browser. Once you have the features installed, you will need to manage them. In the Edge browser you can manage these add-ons by clicking the three dots in the upper-right corner, then choosing Extensions and then Manage Extensions from the drop-down menu. This will open a page where you can see all the installed extensions for the web browser, as shown in Figure 6.34. In this screen you can flip the toggles for any of the extensions installed. These toggles will control if the extension is enabled or disabled.

Extensions work by interacting with the pages rendered. There could be an instance where an extension effects the page being rendered. This menu allows you to configure if the extension is enabled or disabled. A good rule of thumb when having those problems is to turn all the additions off and test your page.

Settings

All extensions have a Details page, as shown in Figure 6.35, where you can view the Description, Permissions, Site Access, option to allow it to operate with inPrivate mode, Source, and a button to remove it. When you enter inPrivate mode, no extensions are enabled by default. This can be controlled on the extensions detail page.


FIGURE 6.34 Edge Extension Settings**FIGURE 6.35** Edge Extension Details

Some extensions have the ability to use their own custom configuration. These extensions will have a link Extension Options, under the View In Microsoft Edge Add-Ons Store link. This link will open a custom web page where you can configure the add-on. Alternately you can click on the puzzle piece icon in the upper right of the web browser, click the three dots next to the add-on, and select Extension Options from the context menu if it provides custom configuration.

Experimental Controls and Features

In most modern web browsers today, there is a hidden menu to control experimental features and control basic features. This is a hidden menu, because clicking around can really cause some issue in your web browsing experience. The Chromium web browsers, such as Chrome and Edge, call this menu Experiments, and it can be accessed at `chrome://flags` or `edge://flags`, respectively, as shown in Figure 6.36.

FIGURE 6.36 Edge Experiments screen



The Experiments screen allows you to turn on experimental features, and control basic features like caching. Typically, unless you know what you are looking to achieve, this hidden menu should be avoided. However, if you wanted to control how cache is handled, or how memory is used, this section is where you would move the levers (so to speak).



Exam Objective 2.10, “Given a scenario, apply security settings on SOHO wireless and wired networks,” requires you to know how to apply security settings for small office, home office (SOHO) networks. You should know how to secure both wireless and wired networks.

Securing a SOHO Network (Wireless)

CompTIA wants administrators of small office, home office (SOHO) networks to be able to secure those networks in ways that protect the data stored on them. This objective looks at the security protection that can be added to a wireless SOHO network, while the following section examines similar procedures for a wired network. This knowledge should complement what you have learned in Volume 1 Chapter 9, “Wireless and SOHO Networking.” The concepts in this section will help you secure wireless and wired SOHO networks.

A wireless network is not and never will be secure. Use wireless only when absolutely necessary. If you must deploy a wireless network, here are some tips to make improvements to wireless security:

- Change the default SSID.
- Disable SSID broadcasts.
- Disable DHCP or use reservations.
- Use MAC filtering.
- Use IP filtering.
- Use the strongest security available on the wireless access point.
- Change the static security keys every two to four weeks.
- When new wireless protection schemes become available (and are reasonably priced), consider migrating to them.
- Limit the user accounts that can use wireless connectivity.
- Use a preauthentication system, such as RADIUS.
- Use remote access filters against client type, protocols used, time, date, user account, content, and so forth.
- Use IPsec tunnels over the wireless links.
- Turn down the signal strength to the minimum needed to support connectivity.
- Seriously consider removing wireless access from your LAN.

Changing Default Usernames and Passwords

In addition to those created with the installation of the operating system(s), default accounts are also often associated with hardware. Wireless access points, routers, and similar devices

often include accounts for interacting with, and administering, those devices. You should always change the passwords associated with those devices and, where possible, change the usernames.

If there are accounts that are not needed, disable or delete them. Make certain that you use strong password policies and protect the passwords with the same security that you use for users and administrators. (In other words, don't write the router's password on an address label and stick it to the bottom of the router.)

Changing the SSID

All radio frequency (RF) signals can be easily intercepted. To intercept 802.11 wireless traffic, all you need is a PC with an appropriate 802.11 card installed. Many networks regularly broadcast their names (known as an *SSID broadcast*) to announce their presence. Simple software on the PC can capture the link traffic in the wireless AP and then process this data to decrypt account and password information.

You should change the SSID—whether or not you choose to disable its broadcast—to keep it from being a value that many outsiders come to know. If you use the same SSID for years, then the number of individuals who have left the company or otherwise learned of its value will only increase. Changing the variable adds one more level of security.

Guest Network Isolation

Most guests in your network never need to connect to the organization's servers and internal systems. When guests connect to your wireless network, it is usually just to get connectivity to the Internet. Therefore, a guest service set identifier (SSID) should be created that isolates guest traffic from production traffic. These guest network SSIDs are usually created by default on consumer wireless devices. On enterprise wireless LAN controllers, the guest network typically needs to be created.

Some considerations for the guest network are what is open to guests, how long they have access, how much bandwidth, the SSID name, and the list goes on, depending on your organization. Guest networks usually don't give totally unrestricted Internet access; certain sensitive ports like TCP port 25 (SMTP) are normally blocked. The length of time they have access is another concern. Generally, a guest is just that, a guest. So, 4 hours, 8 hours, or 24 hours of access seem responsible. You should give this a lot of thought because too short a time will create administrative overhead and too long a window of access allows for abuse of service. If you don't expect guest access to your wireless network, then it should be disabled.



Exam Objective 2.3, "Compare and contrast wireless security protocols and authentication methods," requires you to know wireless security protocols and authentication methods. Be sure to know the differences between the various security protocols.

Setting Encryption

It's important to remember that you should always enable encryption for any wireless network that you administer. Choose the strongest level of encryption you can work with. The following are some wireless protocols that you might encounter when securing wireless:

Open Open security is just that—open with no passphrase or authentication protocol. Open security was originally how all wireless access points (WAPs) were shipped to the customer. Open security still has its uses when used in conjunction with guest wireless access.

Wired Equivalent Privacy Shared passphrases are used with Wired Equivalent Privacy (WEP). WEP provides 64- or 128-bit encryption via the shared passphrase. The passphrase can easily be cracked with tools and is no longer used to secure wireless.

Wi-Fi Protected Access Wi-Fi Protected Access (WPA) was standardized by the Wi-Fi Alliance in 2003 in response to the vulnerabilities in Wired Equivalent Privacy (WEP). WPA uses 256-bit keys versus the 64-bit and 128-bit keys WEP used previously. WPA operates in two modes for security: preshared key (PSK), also called personal mode, and enterprise mode. PSK is the most common mode, because it can easily be implemented with a mutual agreed-upon passphrase. Enterprise mode, also called WPA-802.1X, requires a certificate server infrastructure. Enterprise mode uses the 802.1X protocol, RADIUS, and EAP; it is often used in corporate environments.

WPA introduced many improved security features over WEP, such as message integrity checks (MICs), which detect packets altered in transit. WPA also introduced Temporal Key Integrity Protocol (TKIP), which uses the RC4 algorithm for encryption. TKIP provides per-packet keying to prevent eavesdropping on wireless conversations. However, despite the improvements in security, WPA is considered exploitable and is no longer used for wireless security. A common exploit used against WPA is an attack on the helper protocol of Wi-Fi Protected Setup (WPS). WPS is used for consumer ease of setup and should be turned off for security purposes.

Temporal Key Integrity Protocol (TKIP) TKIP uses the RC4 encryption algorithm protocol as its cipher. TKIP seeds the RC4 algorithm with a key that is derived from the MAC address and initialization vector. TKIP also works in conjunction with message integrity checks (MICs) to check the integrity of messages received by the access point. The MIC protocol, also called Michael, is a 32-bit cyclic redundancy check (CRC). If two CRC MICs fail within 60 seconds, the access point requires TKIP to rekey the RC4 seed value.

Wi-Fi Protected Access 2 (WPA2) WPA2, also known as 802.11i, is the successor to WPA. WPA was deprecated in 2006 when WPA2 became a wireless security standard. Just like WPA, WPA2 operates in both personal mode (PSK) and enterprise mode.

WPA2 uses the Advanced Encryption Standard (AES) algorithm to protect data. AES is more secure than the RC4 algorithm used with TKIP. WPA2 replaced TKIP with Counter

Cipher Mode (CCM) with Block Chaining Message Authentication Code Protocol (CCMP). However, TKIP can be configured as a fallback for WPA backward compatibility. Like WPA, WPA2 is exploitable if the WPS service is enabled. WPS should be turned off for security purposes.

Counter Cipher Mode with Block Chaining Message Authentication Code Protocol (CCMP) Advanced Encryption Standard (AES) uses a 128-bit key to seed the AES encryption and a 128-bit cipher block. The prior ciphered text is used to encrypt the next block of text; this type of cipher is called code block chaining (CBC). CCMP-AES also uses a MIC to check the integrity of wireless data received. If the MIC fails, the CCMP-AES rekeys the session.

Wi-Fi Protected Access 3 (WPA3) This is the successor to WPA2 and became part of the mandatory Wi-Fi Alliance certification process in July 2020. WPA3 introduced 192-bit cryptographic strength in WPA3-Enterprise mode. WPA3 still requires CCMP-128 as a minimum requirement for WPA3-Personal mode. The most significant improvement is the use of the Simultaneous Authentication of Equals (SAE) exchange that replaces the preshared key (PSK) exchange used with WPA and WPA2-Personal mode. SAE was originally introduced with IEEE 802.11s.

Disabling SSID Broadcast

One method of “protecting” the network that is often recommended is to turn off the SSID broadcast. The access point is still there and can still be accessed by those who know about it, but it prevents those who are looking at a list of available networks from finding it. This should be considered a very weak form of security because there are still ways, albeit a bit more complicated, to discover the presence of the access point besides the SSID broadcast.

Wireless MAC Filtering

Most APs offer the ability to turn on MAC filtering, but it is off by default. In the default state, any wireless client that knows of the existence of the AP can join the network. When MAC filtering is used, the administrator compiles a list of the MAC addresses associated with the users' computers and enters them. When a client attempts to connect, an additional check of the MAC address is performed. If the address appears on the list, the client is allowed to join; otherwise, they are forbidden from so doing. On a number of wireless devices, the term *network lock* is used in place of *MAC filtering*, but the two terms are synonymous.



Adding port authentication to MAC filtering takes security for the network down to the switch port level and increases your security exponentially.

Frequencies

The frequencies used with wireless local area networks (WLANs) vary by standard. The two main frequencies used are 2.4 GHz and 5 GHz. The 2.4 GHz frequencies are governed by the industrial, scientific, and medical (ISM) radio bands. The 5 GHz frequencies are governed by the Unlicensed National Information Infrastructure (U-NII) radio band. The 5 GHz frequencies have been expanded upon into the 6 GHz frequency range with the introduction Wi-Fi 6 and Wi-Fi 7. The 6 GHz frequency range is governed by the U-NII radio band.




To learn more about Wi-Fi 6 and Wi-Fi 7 refer to Volume 1 Chapter 9, “Wireless and SOHO Networking.”

2.4 GHz

The 2.4 GHz spectrum is governed by the ISM radio band. The 802.11b/g/n standards operate on 2.4 GHz frequencies. The band consists of 14 channels 22 MHz wide. In North America only the first 11 of the channels can be used for wireless. In Japan all 14 channels can be used, and almost everywhere else in the world the first 13 channels can be used. Only 3 of the 14 channels are considered nonoverlapping, as seen in Figure 6.37. The channels of 1, 6, and 11 are considered prime channels for WLAN because they do not overlap with the other channels in the channel plan.


FIGURE 6.37 The 2.4 GHz channel plan



5 GHz

The 5 GHz frequencies are governed by the Unlicensed National Information Infrastructure (U-NII) radio band. The 802.11 a/n/ac/ax/be standards operate on the 5 GHz frequency spectrum.

As seen in Figure 6.38, the band consists of 25 nonoverlapping channels. In North America the 802.11a standard can function on 12 channels consisting of 36, 40, 44, 48, 52, 56, 60, 64, 149, 153, 157, and 161. Each regulatory domain restricts the number of channels and specific channels for the region.

FIGURE 6.38 The 5 GHz channel plan

In North America, the 802.11ac standard can use 25 of the nonoverlapping channels. In Europe and Japan, the channels are limited to the U-NII 1, U-NII 2, and U-NII 2E list of channels. The 802.11n standard only allowed the first 24 channels in North America, because channel 165 is in the ISM band.

6 GHz

The 6 GHz frequencies are also governed by the U-NII radio band. In recent years the U-NII 5, U-NII 6, and U-NII 7 expanded the 5 GHz frequency range into the 6 GHz range. This expansion meant even more channels for 802.11ax and 802.11be operations. The proposed U-NII 8 looks to even expand into the 7 GHz range in the future.

Speed, Coverage, and Security

Speeds will always be higher on 5 GHz and 6 GHz wireless, such as the 802.11 a/n/ac/ax/be standards. The 802.11ac standard can use a maximum of 25 nonoverlapping channels. However, the 802.11 b/g/n standards are restricted to the 2.4 GHz wireless band with three nonoverlapping channels. The 6 GHz wireless band introduced a whopping 60 nonoverlapping channels. You should always be aware of the airspace around your AP. If you are overlapping on a channel, you should make every effort to change the channel on your AP. By changing the channel, you will effectively increase the speed of the connection for all your users within the cell.

Lower frequencies will go further in distance, yet as you go further away from the AP your speed will suffer. High frequencies will tend to go shorter distances, and when you move further away from the AP, speed will sharply decline. Using higher frequencies allows you to lower power and decrease the changes of signals traveling to a public place.



The 802.11 a/n/ac standards operate at 5 GHz exclusively and 802.11 802.11 b/g/n standards operate at 2.4 GHz exclusively. The latest standard of 802.11ax/be operates at 5 GHz, 6 GHz, and 2.4 GHz. 802.11ax/be will steer the client to the best frequency band, also known as band steering. The concept remains the same: the higher the frequency, the shorter the distance; the lower the frequency, the longer the distance.

Radio Power Levels

Consider the radio power level, since the wireless access point has better transmitting power than most mobile devices. Let's use the analogy of two people standing in a field. One person is using a bullhorn to ask the other person the time, and the other person only has their voice to respond. Although they can hear the request, they will not be heard when they answer because they don't have a bullhorn. To fix this problem, the wireless access point should have its power level adjusted so that the client needs to be closer to receive data or associate with the SSID.

From a security standpoint, power levels should be adjusted so that they do not travel past the interior of the organization's building. If they do, then someone sitting in the parking lot or an adjoining building could attempt to infiltrate the wireless network. On the chance that the signal is actually traveling too far, some access points include power level controls that allow you to reduce the amount of output provided.



A great source of information on RF power values and antennas can be found on the Cisco site:

www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/23231-powervalues-23231.html

Antenna and Access Point Placement

Antenna placement can be crucial in allowing clients to reach the access point. For security reasons, you do not want to overextend the reach of the network so that people can get on to the network from other locations (the parking lot, the building next door, and so on). Balancing security and access is a tricky thing to do.

There isn't any one universal solution to this issue—it depends on the environment in which the access point is placed. As a general rule, the greater the distance the signal must travel, the more it will attenuate; however, you can lose a signal quickly in a short space as well if the building materials reflect or absorb it. You should try to avoid placing access points near metal (which includes appliances) or near the ground. They should be placed in the center of the area to be served and high enough to get around most obstacles.

Wireless site surveys can be performed with specialized software that records the strength of the signal on a map you import of the coverage area. Wireless surveys should be performed in your environment before access point placement to help you determine the best placement. They should also be performed after access point placement for fine-tuning of the location and verification of uniform signal strength.

Assign Static IP Addresses

While DHCP can be a godsend, a SOHO network is small enough that you can get by without it issuing IP addresses to each host. The advantage to assigning the IP addresses

statically is that you can make certain which host is associated with which IP address and then utilize filtering to limit network access to only those hosts.

WPS

WPS (Wi-Fi Protected Setup) can help to secure the network by requiring new machines to do something before they can join the network. This often requires the user to perform an action in order to complete the enrollment process: press a button on the router within a short time period, enter a PIN number, or bring the new device close-by (so that near-field communication can take place). It should be noted that during this brief period the wireless access point is susceptible to an attack. Anyone with the passcode or the ability to guess the passcode can gain access to your wireless network. Therefore, WPS is not used outside of SOHO networks and is not found in corporate networks.

Authentication

When setting up a wireless network, you are extending a wired network to a wireless network. Therefore, you must consider how users will authenticate to the wireless network. The following are various concepts of authentication that can be used with wireless and wired networks (though they are more commonly used with wireless networks):

Single-Factor Authentication As previously mentioned, single-factor authentication is weak authentication. It is based on only one of the following factors:

- Something you know
- Something you have
- Something you are
- Somewhere you are
- Something you do

Setting up a wireless network based on a preshared key limits you to a single authentication factor, which everyone consequently shares. The use of a preshared key is a great example of an authentication factor of something that you know. Unfortunately, a preshared key can be compromised because people must share it, and you only have one preshared key per wireless SSID.

Multifactor Authentication Using single-factor authentication alone allows you to be compromised much more quickly. Therefore, using multifactor authentication is desirable for a higher level of wireless security. *Multifactor authentication* requires two or more single-factor authentication methods to be used to secure the wireless network. A common multifactor authentication implementation is the use of the Extensible Authentication Protocol – Transport Layer Security (EAP-TLS) protocol. EAP-TLS requires that a certificate be installed in the operating system along with a matching password for the computer.

RADIUS *Remote Authentication Dial-In User Service (RADIUS)* was originally proposed as an Internet Engineering Task Force (IETF) standard. It has become a widely adopted industry standard for authenticating users and computers for network systems. RADIUS creates a common authentication system, which allows for centralized authentication and accounting.

The origins of RADIUS are from the original ISP dial-up days, as its acronym describes. Today, RADIUS is commonly used for authentication of virtual private networks (VPNs), wireless systems, and any network system that requires a common authentication system. RADIUS operates as a client-server protocol. The RADIUS server controls authentication, authorization, and accounting (AAA). The RADIUS client can be wireless access points, a VPN, or wired switches. The RADIUS client will communicate with the RADIUS server via UDP port 1812 for authentication and UDP port 1813 for accounting.

The RADIUS server can be installed on many different operating systems, such as Linux and Windows. Microsoft Windows Server includes an installable feature, called the Network Policy Server (NPS), that provides RADIUS functionality.

TACACS+ Terminal Access Controller Access Control System Plus (TACACS+) is a protocol developed by Cisco, from the original dated protocol of TACACS. Although it was developed by Cisco, it was released as an open standard. The protocol is mainly used for authentication of users on routers and switches to allow management access. It is also used to authenticate users connecting to wireless access points via a centralized database. The TACACS+ protocol is declining in popularity and has largely been replaced by RADIUS.

Kerberos Kerberos is an open standard for authentication developed by the Massachusetts Institute of Technology (MIT). It is also the main authentication protocol used by Microsoft Active Directory. Kerberos can be used as the main authentication method for the Microsoft implementation of a RADIUS server, called the Network Policy Server (NPS). Kerberos can also be used with the 802.1X protocol for direct authentication with the Extensible Authentication Protocol (EAP).

Securing a SOHO Network (Wired)

Although a wired network can be more secure than a wireless one, there are still a number of procedures that you should follow to leave as little to chance as possible. Among them, change the default usernames and passwords to different values and secure the physical environment. You should also disable any ports that are not needed, assign static IP addresses, use IP filtering, and use MAC filtering to limit access to hosts that you recognize.

Changing Default Credentials

When installing a network device, the very first thing you must do is log in to the device. There is often a standardized default username and password for each vendor or each vendor's product line. Most devices make you change the default password upon login to the device.

Changing the default password to a complex password is a good start to hardening the device. However, changing the username will also ensure that a brute-force attack cannot be performed against the default username. There are many different websites dedicated to listing the default credentials for network devices, so it doesn't take tremendous skill to obtain the default username and password of a device.

The hosts in the network are no exception to changing default usernames and passwords. In Windows, the Guest account is automatically created with the intent that it is to be used when someone must access a system but lacks a user account on that system. Because the Guest account is so widely known to exist, you should not use this default account but instead create another account for the same purpose if you truly need one. The Guest account leaves a security risk at the workstation and should be disabled to deter anyone attempting to gain unauthorized access.



Change *every* username and password that you can so that they vary from their default settings.

Upgrading Firmware

When you purchase a network device, you don't know how long it's been sitting on the shelf of a warehouse. In that time, several exploits could have been created for vulnerabilities discovered. It is always recommended that a device's *firmware* be upgraded before the device is configured and put into service.

Most hardware vendors will allow downloading of current firmware. However, some vendors require that the device be covered under a maintenance contract before firmware can be downloaded. It is also best practice to read through a vendor's change log to understand the changes that have been made from version to version of firmware.

Filtering

IP filtering, also known as firewall rules, helps secure the internal network from an external network. The external network could be the Internet, or it could be a network less trusted than the internal network, such as a wireless network. In any case, firewall rules help harden the security of an organization because we can restrict activity from the external network to specific applications.

Firewall rules are normally configured with an implicit deny at the end of the rules set. This means that if an application has not explicitly been allowed, it will automatically

(implicitly) be denied. The easy way to remember what *implicit* means is that it implies there is a deny, unless a user or application has been explicitly allowed. This implicit deny operation of a firewall is the default for firewall rule sets.

Although there is an implicit deny at the end of firewall rule sets, you may also need to explicitly deny an application or IP address. An explicit deny is required when another explicit rule follows, allowing access to a wide range of applications or IP addresses. For example, you may want to allow access to all the servers from the client networks. However, you should explicitly deny applications the clients shouldn't have access to, such as Remote Desktop Protocol (RDP) or Secure Shell (SSH).

Disabling Ports

Disable all protocols on the network device that are not required. As an example, many network multifunction printers are preconfigured with a multitude of protocols, such as TCP/IP, Bonjour, and Internet Printing Protocol (IPP), just to name a few. If you don't need them, remove the additional protocols, software, or services, or prevent them (disable them) from loading. Ports on a switch, router, or firewall not in use present an open door for an attacker to enter and should be disabled or disconnected.



Many of the newer SOHO router solutions (and some of the personal firewall solutions on end-user workstations) close down the ICMP ports by default. Keep this in mind, because it can drive you nuts when you are trying to see if a brand-new station, server, or router is up and running.

MAC Address Filtering

Limiting access to the network by employing IP filtering is not the only way to restrict access. *MAC address filtering* can also be employed to restrict traffic to MAC addresses that are known and to filter out those that are not. Even in a home network, you can implement MAC filtering with most routers and typically have an option of choosing to allow only computers with MAC addresses that you list or deny only computers with MAC addresses that you list.



If you don't know a workstation's MAC address, you can find it by using `ipconfig /ALL` (on Windows machines; it is listed as *physical address* and `ifconfig` command on macOS and the `ip` command on Linux operating systems).

Content Filter


Content filters are useful in networks to restrict users from viewing material that is non-work-related, questionable, or malware. Content filtering is usually dictated by organization policy and management. The content filter operates by watching content and requests from

web browsers and other applications. The content filter functions in two ways: the first is content based; when images and text are requested from a website, the content filter can use heuristic rules to filter the content according to administrator-set policies. The second method is URL based, which is much more common since many websites now use SSL/TLS (encryption) and the traffic is encrypted. Content filters are typically purchased with a subscription that provides updates to the categories of the material administrators' block. Content filters can be hardware solutions or software solutions, although it is common to find them installed as software solutions.


Screened Subnet

The *screened subnet* is also known as the demilitarized zone (DMZ). The DMZ gets its name from the segmentation that is created between the exterior and the interior of the network. This is similar to where borders of two opposing countries meet with military presence on both sides. Between the two sides there is a neutral segment called the DMZ. As it pertains to a network, hosts that serve Internet clients are placed in the screened subnet. As shown in Figure 6.39, a network segment called the screened subnet sits between an external firewall and the internal firewall. The external firewall contains ACLs to prevent Internet hosts from accessing nonessential services on the server in the screened subnet. The internal firewall restricts which hosts can talk to internal servers. A typical rule on the external firewall would allow HTTP access for a web server in the screened subnet and would restrict all other ports. A typical rule on the internal firewall would allow only the web server to communicate with the SQL backend database in the internal network.


FIGURE 6.39 A typical screened subnet with two firewalls



The concept of the screened subnet (DMZ) is still used today in network design; a screened subnet can be created between any two segments in the network. The subnets don't necessarily need to be external and internal in relation to the network. Routers containing ACLs can be implemented in lieu of firewalls to filter traffic to the screened subnet, as shown in Figure 6.40. In the figure, a network called Network A is segmented from the screened subnet by a router with ACLs filtering traffic. On the other side of the screened subnet is another network called Network B, and it too is segmented by a router with ACLs filtering traffic. Each of these two networks have equal access to the hosts in the screened subnet. These two networks, Network A and Network B, could potentially be a wireless network and the wired network, respectively.

FIGURE 6.40 A typical screened subnet with two routers

Some screened subnets are just another interface on a single firewall, as shown in Figure 6.41. In this example, the rules for both the Network A subnet and the Network B subnet would be on the same firewall. The benefit of a single firewall is centralized administration of firewall rules. Each interface is placed into a trust zone, and the firewall rules allow incoming and outgoing connections.

FIGURE 6.41 A typical screened subnet with one firewall

Port Forwarding/Mapping

On the router, the port configuration dictates what traffic is allowed to flow through. The router can be configured to enable individual port traffic in, out, or both; when you implement this type of configuration it is referred to as port forwarding. If a port is blocked (such as port 80 for HTTP or port 21 for FTP), the data will not be allowed through, and users will be affected. *Port forwarding* is also known as port mapping, and both are subsets of what a firewall does, and the amount of tweaking they require to get right is about the same.

Port forwarding is required when you are trying to share a network service with the Internet, such as a web server. It is commonly required when gaming, since some games require you to set up an impromptu game server. The players on the Internet will have to connect to the ports on your local machine behind your router, through port forwarding. There are

websites dedicated to configuring these port forwarding settings manually; an example is <https://portforward.com>.

Universal Plug and Play (UPnP)

If you need to configure port forwarding frequently, then you might want to skip manually port forwarding altogether. The *Universal Plug and Play (UPnP)* is a network protocol that allows for automatic configuration of port forwarding. Most modern-day routers have the feature turned on, which is a security concern if you are not using any port forwarding. The UPnP protocol operates by the clients initiating a connection to the router and communicating the ports needing to be forwarded to the client. The router then opens the ports and forwards them to the client.

DHCP

Dynamic Host Configuration Protocol (DHCP) is responsible for automatic configuration of IPv4 IP addresses and subnet masks for hosts from a pool of IPv4 addresses. It is also responsible for configuration of such options as default gateways, DNS server addresses, and many other IP-based servers. It performs configuration of the host in a series of network broadcasts and unicasts.

Reservations

When a client requests an IP address from a DHCP server, the client's MAC address is transmitted in the DHCP packet. A rule on the DHCP server called a *DHCP reservation* can tie the client's MAC address to a particular IP address. When a reservation is created for a client, the client is guaranteed to obtain the same IP address every time for the DHCP process. When a reservation is created on the DHCP server, no other hosts can obtain the reservation IP address unless they have the MAC address that matches the reservation. This type of assignment is considered a dynamically static-assigned IP address.

Reservations can be very handy when static IP addresses are too troublesome to configure, such as network printers with poor configuration options. It's common to set a reservation on network printers and move on when faced with a finicky static IP address process. You can save an hour of busy work in the right situation. Reservations can also be useful when you need to make specific firewall rules for a client based on its IP address.

Static

Dynamic IP addressing is the standard in small-to-large networks when configuring client computers. Static IP addressing should only be used under certain circumstances for client computers, since it is not very scalable and a nightmare to keep track of manually. DHCP allows for central management of the IP address space versus static assignment of individual hosts (which is decentralized). Static IP addressing should only be used on internal network resources such as routers, network printers, and servers.

Static IP addressing can be useful for wide area network (WAN) connections, also known as your connection to the Internet. If a server is operating at the location, a static IP address is necessary for clients to be able to connect. Name resolution to the IP address is the biggest driver for static IP addressing. There are work-arounds, such as dynamic DNS services, but the best solution is to purchase a static IP address from the Internet provider.

Secure Management Access

Configuring secure management access on both wireless and wired routers is critical to protect your network from unauthorized access and potential threats. Start by changing the default login credentials to strong, unique passwords. You should use secure protocols like Secure Shell (SSH) instead of Telnet for remote access. Configuring the router to use HTTPS vs. HTTP for accessing the router's management interface. HTTPS encrypts the data transmitted between your device and the router, protecting sensitive information from being intercepted by malicious actors. Additionally, consider setting up a dedicated management network to isolate management traffic from regular network traffic, further enhancing security.

Some routers offer an option to restrict management access to only the internal network. This means that management interfaces can only be accessed from devices within the local network, adding an extra layer of security by preventing external access.

Physical Security

Just as you would not park your car in a public garage and leave its doors wide open with the key in the ignition, you should educate users to not leave a workstation that they are logged into when they attend meetings, go to lunch, and so forth. They should log out of the workstation or lock it. "Lock when you leave" should be a mantra they become familiar with. A password (usually the same as their user password) should be required to resume working at the workstation.

Also don't overlook the obvious need for physical security. Adding a cable to lock a laptop to a desk prevents someone from picking it up and walking away with a copy of your customer database. Every laptop case we are aware of includes a built-in security slot in which a cable lock can be added to prevent it from easily being carried off the premises. If this is in your home office, it might deter a burglar from walking away with sensitive information.

When it comes to desktop models, adding a lock to the back cover can prevent an intruder with physical access from grabbing the hard drive or damaging the internal components. You should also physically secure network devices, such as routers, access points, and the like. Place them in locked cabinets, if possible—if they are not physically secured, the opportunity exists for an unauthorized person to steal them or manipulate them to connect to the network.



Exam Objective 2.8, "Given a scenario, apply common methods for securing mobile devices," requires you to be able to identify and apply security to mobile devices.

Mobile Device Security

Apple computers have a pretty decent reputation in the industry for being somewhat resistant to malware. Whether or not this is because of the relatively small installed base or the ease with which hackers penetrate “other” operating systems, this characteristic carries over to Apple’s mobile devices. In fact, hackers don’t seem to be as interested in attacking the legions of mobile devices as much as they have gone after the Windows operating systems that drive the vast majority of laptops, desktops, and servers in the world. Nevertheless, attacks occur. Coupled with how easy mobile devices are to misplace or steal, it behooves users to have proactive monitoring and contingency plans in place.

The following sections detail the built-in security utilities that are common in today’s mobile devices. Furthermore, for threats not covered by the software with which the devices ship, the protection available from third-party utilities is worth discussing.

Hardening Techniques

Hardening techniques for mobile devices enhances security and protecting sensitive data from potential threats and threat actors. By implementing these techniques, users can significantly reduce the risk of unauthorized access, data breaches, and other cyberattacks. In the following sections, you will learn about various hardening methods, including screen locks, which prevent unauthorized access to the device; device encryption, which ensures that data stored on the device is secure; endpoint detection and response (EDR) software, which protect against malicious software; firewalls, which help control incoming and outgoing network traffic; and configuration profiles, which enforce security policies and settings. These measures collectively contribute to a robust security posture, safeguarding mobile devices against a wide range of threats.

Screen Locks

Apple and Android mobile devices include a requisite locking mechanism, which is off by default. The user on the go is encouraged to enable a lock. If your device acts primarily as a home computing device and rarely goes with you out the door, there is very little reason to set a lock. However, knowing how to do so is important. The following are types of locks that you can implement to secure your device:

Fingerprint Lock The *fingerprint lock* is a biometric-type lock that uses your fingerprint to unlock the device. The fingerprint lock is the most secure biometric method for a device. The technology works by placing your finger on a touchpad sensor on the device. On Apple devices, the fingerprint recognition is called *Touch ID*. The technology works exceptionally well, except when your fingers are wet or extremely dirty. When your fingers are wet or extremely dirty, they make different prints from those you originally trained the device to recognize.

Face Lock or Face ID The *face lock* and *face ID* are biometric-type locks that allow you to use your face to unlock your device. There are a high number of false positives that makes the face lock less secure than the fingerprint lock. This means that presumably someone that looks like you could unlock your phone. Almost all mobile devices have a front-facing camera, allowing you to instantly know that you've unlocked the phone seamlessly by looking at your phone. On Apple devices, this technology is called *Face ID*.

Pattern Lock The *pattern lock* is a type of passcode that allows you to unlock your device. Swipe lock works by displaying nine dots in a matrix of 3×3 . You then swipe with your finger with the registered pattern to unlock the phone. The swipe lock is the least secure of any locking methods. A grease trail from your fingers can allow someone to derive the swipe pattern.

PIN Lock The *PIN lock* is by far the most common way to unlock devices. It is typically a 4- or 6 digit-numeric passcode. However, it can also be alphanumeric, depending on the device. Passcode locks suffer from the same problems that passwords are prone to: people can shoulder surf or learn your passcode over time.

Swipe Lock The *swipe lock* is not really a lock at all. It allows you to unlock the phone easily by swiping on the screen. This lock method is to really prevent pocket dialing of someone by accident.

Exercise 6.4 outlines the steps for creating a code for your iPhone.

EXERCISE 6.4

Setting the Passcode Lock on an iPhone

1. Tap Settings.
2. Depending on the iOS version, select either Face ID or Passcode.
3. Tap Turn Passcode On.
4. Enter a six-digit passcode.

If you want to use a four-digit passcode, tap Passcode Options to switch to a four-digit numeric code or a custom alphanumeric passcode.

5. Enter the passcode a second time to confirm. Doing so sets the passcode.

Optionally, you can set the amount of time that must pass while the phone is asleep before the passcode will be required and whether the Erase Data feature should be enabled. Setting the Required Passcode field to Immediately requires entering the passcode each time the device wakes up.

The same general concept for Android phones is illustrated in Exercise 6.5.

EXERCISE 6.5

Setting the Passcode Lock on an Android Phone

Depending on the Android version, these steps might differ.

1. Swipe up on the home screen.
 2. Tap the Settings app.
 3. Select Lock Screen and/or Security from the Settings menu.
 4. Tap Screen Lock Type on the Lock Screen And Security menu.
 5. Select Pattern from the Phone Security list.
 6. Use your finger to draw a continuous pattern of four or more dots, and then tap the Continue button.
 7. Repeat the same pattern, and then tap the Confirm button.
 8. Optionally, tap Secure Lock Settings to lock the phone automatically and adjust the number of wrong swipes before a factory reset occurs.
-

Full-Device Encryption

As discussed previously, BitLocker and BitLocker to Go greatly enhance security by encrypting the data on drives (installed and removable, respectively) and helping to secure it from prying eyes. At a minimum, the same level of protection that you would apply to a desktop machine should be applied to a mobile device, because it can contain confidential and personally identifiable information (PII), which could cause great harm in the wrong hands.

Full-device encryption should be done on laptops and mobile devices, and you should back up regularly to be able to access a version of your files should something happen to the device itself. When full-device encryption is turned on, both the device and the external storage (SD card) are encrypted. The only way to view the information on the SD card is from the encrypting mobile device.



Android devices do not automatically turn on device encryption and encrypt files upon setting a passcode on the device, like iOS devices do.

Failed Login Attempt Restrictions

After you've set a screen lock, an optional step is to set the device to wipe, or factory reset, after a number of failed attempts. This option will wipe local data on the device if incorrect

passcodes are entered 10 times in a row, or perform a factory reset depending on the device. While this is recommended for users with devices that contain sensitive data and that are frequently taken into public venues or placed in compromising positions, the casual user should not turn this feature on unless they can be sure that a recent backup will always be available in iCloud or Google.

Imagine a user's child or a mischievous, yet harmless, friend poking away at passcodes until the device informs them that it is being wiped clean; it's not for everyone. Restoring from a backup is easy enough, but will a recent backup be available when disaster strikes? Apple performs a backup to the iCloud or the computer running iTunes that the iOS device syncs with.

Apple imposes cooling-off timeout periods of increasing duration, even if the Erase Data feature is disabled and you or someone else repeatedly enters the wrong code over multiple lockouts. The final penalty with the Erase Data feature disabled is that you cannot unlock the device until it is connected to the computer with which it was last synced.

When a passcode is set, Android devices have a similar approach to failed login attempts as their Apple counterparts. A factory reset will occur after 10 failed attempts.

The difference is that if waiting the timeout period won't help because you've forgotten the pattern or code, this device can tie your access back to the Google account you used when setting it up. This is also the account where you receive purchase notifications from Google Play, and it does not have to be a Gmail account (one of the benefits of the open-source nature of Android). If you remember how to log into that account, you can still get into your phone. Or at least you can investigate the credentials to that account on a standard computer and return to the device to enter those credentials.

Endpoint Detection and Response (EDR)

For the most part, mobile devices have been left alone with viruses and malware, as compared with the Windows platform. However, this is not a reason to let your guard down and not worry about viruses and malware on mobile devices. Mobile devices can contract a virus and malware through the installation of a malicious app.

Most malware installed on mobile devices will spam your device with ads from the Internet. However, some malware can expose your personal information and even subject your device to the control of a malicious individual. Viruses, on the other hand, will use your email application to send copies of itself or turn your device into a zombie.

The extent of the damage from malware or a virus is hard to estimate, but one thing is for sure: protection goes a long way. Antivirus and antimalware software should be installed on your mobile device to thwart malicious attempts to infiltrate your device. All the leading vendors, such as AVG, Norton, and Avast (just to name a few), have offerings of antivirus and antimalware. Most installations of these apps are free, and other services or features can be purchased within the apps.

Antivirus software primarily targets known threats such as viruses, worms, and trojans, while antimalware focuses on a broader range of malicious activities, including spyware, ransomware, and adware. Endpoint Detection and Response (EDR) for mobile devices integrates these capabilities to offer advanced threat detection and response. EDR continuously monitors mobile devices for suspicious activities, collects and analyzes data, and provides real-time alerts

to security teams. This enables rapid identification and mitigation of potential threats, reducing the risk of data breaches and ensuring the security of sensitive information on mobile devices.

EDR solutions for mobile devices go beyond traditional antivirus and antimalware by incorporating advanced features such as behavioral analysis, machine learning, and threat intelligence. These technologies enable EDR to detect and respond to sophisticated threats that may evade conventional security measures.

Content Filtering

Mobile content filtering is a technology used in environments such as schools, businesses, and government organizations where it is essential to ensure that users are not exposed to harmful or inappropriate content. By using mobile content filtering, administrators can block access to specific websites, apps, and online services that may pose security risks or distract users from their tasks. Mobile content filtering helps protect against malware infections and phishing attacks by blocking access to malicious websites and links. This is particularly important for maintaining the security of organizational data and ensuring that employees stay on task by preventing access to irrelevant sites that reduce productivity and drain bandwidth. Additionally, mobile content filtering can be configured to comply with legal and regulatory requirements, ensuring that organizations adhere to data protection and privacy laws.

Mobile content filtering works by screening content to restrict or block access to websites and online services that compromise security or breach internal content consumption policies. This is achieved through various techniques such as allowlists and blocklists, keyword filtering, and AI-powered filters that determine allowable data. For instance, allowlists include a small number of allowed sites and block access to everything else, while blocklists apply the opposite principle. AI-powered filters add a valuable layer of subtlety by adapting to different contexts and blocking harmful content without affecting legitimate sites.

Trusted Sources vs. Untrusted Sources

You can limit your exposure to malicious apps by only installing apps from trusted sources, such as the Apple's App Store or Google Play. In most cases, the mobile operating system must specifically be configured to accept installations from untrusted sources. So, it's not likely that you will mistakenly install an app from an untrusted source. Examples of untrusted sources include manual installs of Android APK (Android Package Kit) or untrusted Apple IPA (iOS App Store Package) files. These files can be distributed outside of the Google Play and Apple's App Store ecosystems. When files are distributed outside of these ecosystems, consider them untrusted.

Firewalls

Mobile devices have the same inherent vulnerabilities as desktop operating systems. However, mobile devices are generally not targeted the same as desktop operating systems. When a mobile device starts up, it is configured with an IP address via either the internal cellular radio or the Wi-Fi radio, so the device can be exploited via the network. Therefore, a firewall should be installed and turned on to protect the mobile device.

However, inbound communication is not the only aspect of security you need to be concerned about on mobile devices. Outbound communication also is a primary concern. A mobile device firewall app will allow you to monitor both the inbound and outbound communications on your mobile device. Several third-party firewall apps (both free and for purchase) can be downloaded from either Apple's App Store or Google Play.

Configuration Profiles

Configuration profiles are a key component of mobile device management (MDM). They contribute a major role in maintaining the security and compliance of devices within an organization. These profiles are essentially sets of rules and settings that administrators can deploy to mobile devices to control various aspects of their operation. By using configuration profiles, organizations can ensure that all devices adhere to the same security standards and policies, which is vital for protecting sensitive data and preventing unauthorized access.

For example, Microsoft Intune, a cloud-based MDM solution, allows administrators to create and deploy configuration profiles to manage mobile devices. With Intune, you can configure a wide range of settings, including:

Security Policies: Enforce password requirements, such as complexity and expiration, to ensure that all devices have strong passwords.

Encryption: Require device encryption to protect data at rest, ensuring that even if a device is lost or stolen, the data remains secure.

App Management: Restrict access to certain apps or prevent the installation of unauthorized apps, reducing the risk of malware and other security threats.

Network Settings: Configure Wi-Fi, VPN, and other network settings to ensure secure and reliable connectivity.

For instance, an organization might use Intune to create a configuration profile that enforces encryption on all corporate devices. This profile would automatically apply the encryption settings to any device enrolled in Intune. This would ensure that all data stored on the device is encrypted and protected. Additionally, the profile could include settings to require a complex passcode, further enhancing the security of the device.

Remote Wipes and Locator Applications

Should your work or personal mobile device disappear or fall into the wrong hands, it's always nice to have a backup plan to ensure that no company secrets or personal identifiers get misused by anyone who would use the information with ill will. Apple supplies a free app called *Find My iPhone* that, together with iCloud, allows multiple mobile devices and Macs to be located if powered on and connected to the Internet (via cellular, Wi-Fi, Ethernet,

and so on). The app allows the device to be controlled remotely to lock it, play a sound (even if audio is off), display a message, or wipe it clean.

Within a newer iPhone's Settings screen, you can find an iCloud settings page and select the Find My iPhone switch. If this switch is off, the Find My iPhone app and iCloud web page will be unable to find your device.

On the login screen for the iPhone app, you must enter the iCloud account information that includes the device you are attempting to control remotely.

Note that when you change the password for your Apple ID, it affects your iCloud account but does not transfer automatically within your device. Always remember to update your iCloud account information on each device when you update the associated Apple ID.

The iCloud website's login page (www.icloud.com) calls for the same credentials as the app. You are signing in with HTTPS, so your username and password are not traversing the Internet in the clear. With the switch in the iCloud settings screen set to Off for all devices on your account, when you sign in to the app with your iCloud account credentials, you are met with a disabling switch message.

You do not need to go to the website if you have another device with the Find My iPhone app or can borrow one from someone else. The device forgets your credentials when you log out, so the owner will not be able to control your device the next time they use the app.

After logging into the iCloud website, you can click the icon that matches the icon for the Find My iPhone app in iOS. Assuming that you've made it into the app on another device and your Find My iPhone feature is enabled on your missing device, the Info screen tells you that your device has been found and gives you options for the next step you take.

Tapping the Location button in the upper left shows you a map of where your device is currently located. You have three options for how to view the location: a two-dimensional map, a satellite map, and a hybrid version, where the two-dimensional street-name information is laid over the satellite view.

If you tap the Play Sound or Send Message button on the Info screen, instead of the Location button, the screen that pops up allows you to display a message remotely. You might consider first displaying a message without the sound, which is at maximum volume. Ask in the message to be called at another number. If you hear from someone in possession of your device, the hunt is over. Otherwise, send another message with the tone to get the attention of the nearest person. If you are at the reported location when you generate the sound, it can help you home in on the device.

If you do decide to use the remote-lock feature of the app, you'll have the opportunity to reconsider before locking it. You should have no issue with locking the device. Doing so does not prevent you from using the app further; it simply makes sure that the device is harder to break into.

Should you decide to take the sobering step of destroying the contents of the device remotely, you get a solemn notice allowing you the opportunity to reassess the situation before proceeding. Sometimes there's just no other option.

For Android devices, Google's *Find My Device* app performs many of the same functions as Apple's Find My iPhone app, including playing your ring tone for 5 minutes. When you

search Google's web page for "find my device," a map will display the location of your device. Clicking the location will reveal the SSID the phone is connected to, the phone's battery level, and when the phone was last in use at that location. You can then choose to play a sound (ring tone), secure the device, or erase the device. Securing the device will lock the phone and sign you out of your Google account. This option can also allow you to display a message on the lock screen. Erasing the device is the last step when you know it cannot be recovered and you want to be assured your data is removed.



Real World Scenario

Uh-Oh. Where's My Phone?

After a recent visit to a car dealership, a prospective new car owner named Jimmy departed without his iPhone—at least, that's where Find My iPhone said it was. Luckily, Jimmy had the foresight to set up his free iCloud account and log that phone into it. Upon his return to the dealership, Jimmy was temporarily disheartened to find that no one had seen his phone.

It occurred to him that he had laid his phone down on the counter in the men's room as he was washing his hands. He didn't recall having it after that point. A cursory look around the facilities turned up nothing. The app was not accurate enough to tell him where exactly the phone was at the time it was located on the Find My iPhone website, which he checked from home. However, the app did indicate that the iPhone was still somewhere at the dealership.

Jimmy noticed that the salesperson assisting him had an iPhone, so he asked her if she also used the Find My iPhone app. She did. He asked if he could borrow her phone for a brief instant. She obliged, and Jimmy entered his iCloud credentials into the salesperson's app and then entered the message, "I'm in the showroom." He left the Play Sound switch on and tapped the Send button.

In less than five minutes, a manager who had been back at the loading dock came to the showroom with a story. He said he heard a disturbing noise coming from the dumpster, which was quite full and scheduled to be picked up the next morning. The trash bag containing Jimmy's phone was conveniently right at the top of the heap and easily retrieved by the manager.

Jimmy recalled being in a bit of a fluster as he left the men's room. He was about to finalize the terms of the deal he had been working on for more than an hour. In his haste to get back to the table, apparently his phone slipped from his hand as he was disposing of the paper towels he used to dry his hands. So, the cliché "All's well that ends well" has a shred of truth. Jimmy owed this happy ending to the features of the Find My iPhone app and tells his story to anyone who will listen.

Remote Backup

Apple iOS devices automatically back themselves up either to a computer running iTunes that they sync with or to the iCloud account associated with the device. When a mobile device is connected to a computer containing iTunes, the process of backup is called *synchronization*. iCloud Backup is the most common method of backing up Apple iOS devices, since no computer is necessary. iCloud does require a Wi-Fi connection or cell service. To enable iCloud Backup, navigate to Settings, then tap your name, tap iCloud, and finally tap iCloud Backup. You can create an immediate backup by tapping Back Up Now, or you can wait until the next backup interval.

The Android operating system will automatically synchronize the device to Google Drive. Android phones require a Google account during setup. This account is the account that will synchronize to Google Drive. Google Drive is Google's cloud-based storage product. The backup service will back up Wi-Fi passwords, phone logs, app settings, contacts, messages, pictures, and other related files. A multitude of third-party backup apps can be downloaded from Google's Play Store. Each third-party app offers different features over and above the built-in backup service functionality, such as the capability to remove the bloatware apps that come with the phone.



Both Apple iCloud and Google Drive services allot you a certain amount of storage when you sign up for free. The limited cloud storage used for backups of your settings and data is adequate for most people. However, you can always purchase more storage from Apple and Google.

Patch Management

Regularly applying patches helps to fix vulnerabilities that could be exploited by threat actors, thereby protecting sensitive data and preventing unauthorized access. For mobile devices, this means not only updating the operating system but also ensuring that all installed applications receive the latest security patches and feature enhancements. Effective patch management reduces the risk of malware infections, improves device performance, and ensures compliance with regulatory requirements. By keeping both the operating system and applications current, organizations can safeguard their mobile environments against emerging threats and keep them secure.

Operating System Updates

It's easy to forget that these tiny yet powerful mobile devices we've been talking about are running operating systems that play the same role as the operating systems of their larger siblings. As such, users must be careful not to let the operating systems go too long without updates. Occasionally, mobile devices will notify the user of an important update to the operating system. Too often, however, these notifications never come. Therefore, users should develop the habit of checking for updates on a regular basis.

Not keeping up with software updates creates an environment of known weaknesses and unfixed bugs. Mobile devices operate on a very tight tolerance of hardware and software performance. Not maintaining the device for performance at the top of its game will tend to have more pronounced repercussions than those seen in larger systems.

For the iPhone and iPad, you can check for the most important level of updates by tapping Settings > General > Software Update. For the Android operating system, there are multiple updates that can be checked for manually. All of them are accessible by following Menu > Settings > System Updates.

Application Updates

On Android devices, users can enable auto-updates by navigating to the Google Play Store settings and selecting their preferred update method, either over cellular data or Wi-Fi only. This setting allows the device to automatically check for and install updates, ensuring that apps are always current without user intervention. Similarly, Apple devices offer an auto-update feature that can be configured for both cellular data and Wi-Fi only downloads. Navigating to Settings > Apps > App Store, you can configure the application auto-updates. By default, both Android and Apple devices will auto-update applications over Wi-Fi, but not over cellular data.

When auto-updates are enabled on Android devices, the Google Play Store periodically checks for updates and installs them in the background. This process is designed to be seamless, minimizing disruption to the user. On Apple devices, the App Store handles automatic updates similarly. Users can also manually update apps by opening the App Store, tapping on their profile icon, and selecting Update All or updating individual apps as needed.

Multifactor Authentication

Multifactor authentication, mentioned previously in this chapter, involves using more than one item (factor) to authenticate. An example of this would be configuring a BitLocker to a Go-encrypted flash drive so that when it is inserted into your laptop, a password and smart-card value must be given before the data is decrypted and available.

Authenticator Applications

An authenticator app works with mobile devices to generate security codes that can keep accounts secure by requiring two-factor authentication (2FA). Once this is set up, your account will require a code from the app in addition to your account password. Authenticator applications are available for download with Apple, Android, and other mobile device operating systems, as well as desktop operating systems. An account is usually added to the authenticator application by entering a secret key or scanning a QR barcode; this creates the account in the authenticator application. Several different authenticator apps can be downloaded; they vary depending on the mobile operating system.

Policies and Procedures

With the explosive growth of mobile devices in the workplace, there are many different policies and procedures that may be required for your organization to minimize data loss. This section focuses on the policies and procedures specific to mobile devices that you may have in your organization.

As employees are hired in your organization, a certain amount of initial interaction with the information technology (IT) is required. This interaction, called the *onboarding procedure*, is often coordinated with the Human Resources (HR) department. During the onboarding procedure, the IT representative will help the user log in for the first time and set their password. The password policy is often the first policy discussed with the user. Other policies—such as bring your own device (BYOD), acceptable use policy (AUP), nondisclosure agreement (NDA), and information assurance—should also be discussed during the onboarding procedure. The use of email and file storage and policies should also be covered with the user during the onboarding process. Each organization will have a different set of criteria that make up the onboarding procedures.

Eventually, employees will leave your organization. The *offboarding procedure* ensures that information access is terminated when the employment of the user is terminated. The offboarding procedure will be initiated by the HR department and should be immediately performed by the IT department. This process can be automated via the organization's employee-management system. This procedure can also be performed manually if the employee-management system is not automated. However, the procedure must be performed promptly, since the access to the company's information systems is the responsibility of the IT department.

During the offboarding procedure, email access is removed via the *mobile device management (MDM)* software, the user account is disabled, and IT should make sure that the user is not connected to the IT systems remotely. The offboarding procedure may also specify that the user assume ownership of the terminated employee's voicemail, email, and files.

BYOD vs. Corporate-Owned

The traditional workforce is becoming a mobile workforce, with employees working from home, on the go, and in the office. These employees use laptops, tablets, and smartphones to connect their companies' cloud resources. Organizations have embraced BYOD initiatives as a strategy to alleviate the capital expense of equipment by allowing employees to use devices they already own.

Because employees are supplying their own devices, a formal document called the *BYOD policy* should be drafted. The BYOD policy defines a set of minimum requirements for the devices, such as size and type, operating system, connectivity, antivirus, patches, and many other requirements the organization will deem necessary, as well as the level of service your IT department will support for personally owned equipment.

Many organizations use MDM software, which helps enforce the requirements for the BYOD policy. MDM software helps organizations protect their data on devices that are

personally owned by the employees. When employees are terminated or a device is lost, the MDM software allows a secure remote wipe of the company's data on the device. The MDM software can also set policies requiring passwords on the device. All of these requirements should be defined in the organization's BYOD policy.

Corporate-owned mobile devices are also of paramount concern when it comes to security. The equipment is mobile, so these devices sometimes travel and can disappear completely. Luckily, MDM software allows you to not only control the data on these devices, but in many instances even track it via a built-in Global Positioning System (GPS) sensor (if the device supports this functionality). When MDM software is implemented in the capacity of tracing assets, consideration must be given to privacy. Although it may be acceptable to track a corporate-owned mobile device, the end user must be made aware of this policy.

Profile Security Requirements

When you implement an MDM solution to manage mobile devices, part of your implementation is the creation of profile security requirements for the mobile devices you will manage. The profile security requirements allow the management of the mobile devices in a uniform fashion. As an administrator, you can choose settings for mobile devices under your purview and enforce profile security requirements in various ways. In a given scenario, you may want to enforce settings for the entire organization, whereas in other scenarios you may want to differ the settings based on organizational unit, role, or other group type. Among the settings you may want to enforce are those requiring the encryption of drives and the use of complex passwords.

IoT Considerations

The *Internet of Things (IoT)* is an exploding industry. You can now control everything from your lights to your thermostat with a mobile device or computer. The downside to connecting things to the Internet is that they must be patched so that they are not exploitable. Hardware vendors often wire off-the-shelf components into the hardware, and these off-the-shelf components never get frequent updates. So, this is a major security consideration for an organization. Ease of use versus security is the balance beam that you walk when owning IoT devices.

In recent years, attackers have harvested IoT devices for distributed denial-of-service (DDoS) attacks. The Mirai botnet is one such popular botnet that can crush an organization's bandwidth. This can happen when IoT devices in your network are used in a DDoS or are attacked by the botnet. To mitigate inadvertently being used in a DDoS, you can place the IoT devices on an isolated network and police their outbound bandwidth.

Unfortunately, there is not much you can do with IoT devices to prevent being attacked. If an attacker wants to attack your organization with IoT devices, firewalls, ISP controls, and third-party services like Cloudflare can help mitigate these attacks. This is not really an IoT consideration because any botnet can attack you and the mitigations are the same as if you were attacked by an IoT botnet.

Summary

In this chapter, you learned some best practices related to operating system security. We then focused on Windows operating system security settings, such as sensitive accounts and file-system permissions. We concluded with mobile device security for both Apple devices and Google Android. The mobile device operating systems are changing rapidly, as both consumer and enterprise needs change. However, the objectives on the A+ 220–1202 exam are general enough that a mastery of them will enable you to secure a mobile device today and into the future.

Security, as you’ve already guessed, is a large part of the CompTIA A+ certification. CompTIA expects everyone who is A+ certified to understand security-related best practices and be able to secure both Windows operating systems and mobile devices.

Exam Essentials

Understand Windows operating system security concepts. You should be able to compare and contrast Microsoft Windows operating system security settings. This should include knowing the various accounts that are created by default and why you need to disable or rename these accounts. You should also have an in-depth understanding of NTFS permissions and share permissions, and their interaction, to form the effective permissions the user has when accessing a share remotely. In addition, you should have a good understanding of authentication and the options to encrypt both the operating system and removable drives.

Know how to compare wireless security protocols and authentication methods. Wireless networks can be encrypted through WEP, WPA, WPA2, and WPA3 technologies. Wireless controllers use service-set identifiers (SSIDs)—32-character, case-sensitive strings—that must be configured in the network cards to allow communications. However, using ID string configurations doesn’t necessarily prevent wireless networks from being monitored, and there are vulnerabilities specific to wireless devices.

Understand the basics of antenna placement and radio power levels. Antenna placement can be crucial in allowing clients to reach an access point. Place access points near the center of the area to be served and high enough to get around most obstacles. Know that power level controls allow you to reduce the amount of output provided.

Know how to configure security on SOHO wireless and wired networks in various scenarios. Disable all unneeded protocols/ports. If you don’t need them, remove them or prevent them from loading. Ports not in use present an open door for an attacker to enter. MAC filtering allows you to limit access to a network to MAC addresses that are known and filter out (deny access to) those that are not.

Understand mobile device security concepts. You should understand how screen locks operate. Explain why we need remote wipe in the event of a device missing or lost. Understand the risks of viruses and malware on mobile devices, as well as measures that can be taken to mitigate these risks. In addition, you should have a good understanding of the importance of full-device encryption, authentication applications, firewalls, and policies for mobile devices.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which policy would you create to define the minimum specification if an employee wanted to use their own device for email?
 - A. MDM
 - B. AUP
 - C. BYOD
 - D. NDA
2. Which term refers to copying data between a mobile device and a computer system in order to mirror such things as contacts, programs, pictures, and music?
 - A. Calibration
 - B. Remote wipe
 - C. Pairing
 - D. Synchronization
3. You want to follow the rules of good security administration as set by CompTIA and vendors. To do so, which account should be disabled on most Windows operating systems for security reasons?
 - A. Guest
 - B. Print Operators
 - C. Power Users
 - D. Userone
4. What kind of mobile app is being used when the owner's phone displays a message on the screen and emits an extremely loud tone?
 - A. Failed login restriction
 - B. Antivirus
 - C. Locator
 - D. Remote wipe
5. As a best practice, after a set period of inactivity on a Windows workstation, what should happen?
 - A. The system should shut down.
 - B. The system should restart.
 - C. A password-enabled screensaver should automatically start.
 - D. The system should log out the user.

6. Which Windows feature helps protect against unauthorized changes to system files?
 - A. BitLocker
 - B. EFS
 - C. UAC
 - D. Windows Defender
7. Which of the following has the goal of allowing a username/password combination to be entered once and then allowing claims to be used for consecutive logins? (Choose the best answer.)
 - A. Tokens
 - B. Kerberos
 - C. Single sign-on
 - D. Multifactor authentication
8. Which of these is a password manager?
 - A. Edge
 - B. Credential Manager
 - C. Internet Explorer 11
 - D. Active Directory
9. You have a very small network in a home-based office, and you want to limit network access to only those hosts that you physically own. What should you utilize to make this possible?
 - A. Static IP addresses
 - B. Disabled DNS
 - C. Default subnet mask
 - D. Empty default gateway
10. Which wireless encryption protocol provides Advanced Encryption Standard (AES) encryption?
 - A. Wired Equivalent Privacy (WEP)
 - B. Wi-Fi Protected Access (WPA)
 - C. Wi-Fi Protected Access 2 (WPA2)
 - D. Temporal Key Integrity Protocol (TKIP)
11. Which type of add-on will extend the functionality of the web browser in a way it wasn't originally designed?
 - A. Pop-up blocker
 - B. Extensions
 - C. Plug-in
 - D. Ad blocker

12. What is normally performed when an employee is offboarded?
 - A. Their user account is deleted.
 - B. Their user account is unlocked.
 - C. Their user account is created.
 - D. Their user account's password is reset.
13. By default, when setting up an Android device, why do you need a Google account?
 - A. The device requires email setup.
 - B. The account is used for cloud synchronizations.
 - C. The account is used for desktop backups.
 - D. The device requires registration.
14. You need to secure your mobile device's lock screen with the highest level of protection. Which of the following should you use? (Choose the best answer.)
 - A. Fingerprint lock
 - B. Face lock
 - C. Passcode lock
 - D. Swipe lock
15. You need to encrypt a single file on a Windows desktop. Which technology should you use?
 - A. EFS
 - B. BitLocker
 - C. NTFS
 - D. BitLocker to Go
16. A user is in both the Sales group and the Marketing group. The Sales group has full permission at the share level, and the Marketing group has Read-only permission. The files on NTFS are secured with the Modify permission for the Sales group and the Read & Execute permission for the Marketing group. Which permissions will the user have?
 - A. Full
 - B. Modify
 - C. Read-only
 - D. Read & Execute
17. James just moved a folder on the same partition. What will happen with the permissions for the folder?
 - A. The permissions will be the same as they were before the move.
 - B. The permissions will be inherited from the new parent folder.
 - C. The permissions will be configured as the root folder for the drive letter.
 - D. The permissions will be blank until configured.

- 18.** A user is in the Sales group. The Sales group has no permissions at the share level. The files on NTFS are secured with the Modify permission for the Sales group. What permissions will the user have?
- A.** The user will have the Modify permission when connecting from the network.
 - B.** The user will have the Modify permission when logged in locally to the computer.
 - C.** The user will have no access when logged in locally to the computer.
 - D.** The user will have Read-only permissions when connecting from the network.
- 19.** You are trying to delete a file on the local filesystem, but the operating system will not let you. What could be the problem? (Choose the best answer.)
- A.** The NTFS Modify permission is applied to the file.
 - B.** The share permissions are not set to Full Control.
 - C.** The file attributes are set to Read-only.
 - D.** The file attributes are set to System.
- 20.** You need to enforce profile security requirements on mobile devices. Which should you use to achieve this goal?
- A.** AUP
 - B.** NDA
 - C.** BYOD
 - D.** MDM

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

You have been asked to create a working structure for your organization's network. You have three groups: Sales, Marketing, and R&D. You need to set up a network share and NTFS to allow Sales to access Marketing material but not modify it in any way. R&D must be able to write to marketing files and read Sales information. Marketing must only have read access to R&D and Sales. Each group should have the Modify permission to their respective folder. All permissions should be controlled with share permissions. How will you set up the folders for access, NTFS permissions, and share permissions?



Chapter 7

Troubleshooting Operating Systems and Security

**THE FOLLOWING COMPTIA A+ 220-1202
EXAM OBJECTIVES ARE COVERED IN THIS
CHAPTER:**

✓ **2.0 Security**

✓ **2.6 Given a scenario, implement procedures for basic small office/home office (SOHO) malware removal.**

- Investigate and verify malware symptoms.
- Quarantine the infected systems.
- Disable System Restore in Windows Home.
- Remediate infected systems.
 - Update anti-malware software.
 - Scanning and removal techniques (e.g., safe mode, preinstallation environment).
 - Reimage/reinstall.
- Schedule scans and run updates.
- Enable System Restore and create a restore point in Windows Home.
- Educate the end user

✓ **3.0 Software Troubleshooting**

✓ **3.1 Given a scenario, troubleshoot common Windows OS issues.**

- Blue screen of death (BSOD)
- Degraded performance
- Boot issues
- Frequent shutdowns
- Services not starting



- Applications crashing
- Low memory warnings
- USB controller resource warnings
- System instability
- No OS found
- Slow profile load
- Time drift

✓ **3.2 Given a scenario, troubleshoot common mobile OS and application issues.**

- Application fails to launch
- Application fails to close/crashes
- Application fails to update
- Application fails to install
- Slow to respond
- OS fails to update
- Battery life issues
- Randomly reboots
- Connectivity issues
 - Bluetooth
 - Wi-Fi
 - Near-field communication (NFC)
- Screen does not autorotate

✓ **3.3 Given a scenario, troubleshoot common mobile OS and application security issues.**

- Security concerns
 - Application source/unofficial application stores
 - Developer mode
 - Root access/jailbreak
 - Unauthorized/malicious application
 - Application spoofing
- Common symptoms



- High network traffic
- Degraded response time
- Data-usage limit notification
- Limited internet connectivity
- No internet connectivity
- High number of ads
- Fake security warnings
- Unexpected application behavior
- Leaked personal files/data

✓ **3.4 Given a scenario, troubleshoot common personal computer (PC) security issues.**

- Common symptoms
 - Unable to access the network
 - Desktop alerts
 - False alerts regarding antivirus protection
 - Altered system or personal files
 - Missing/renamed files
 - Inability to access files
 - Unwanted notifications within the OS
 - OS update failures
- Browser-related symptoms
 - Random/frequent pop-ups
 - Certificate warnings
 - Redirection
 - Degraded browser performance



Troubleshooting is a major responsibility of an A+ technician's daily job. It may not be as glamorous as we'd like it to be, but it does make up a good percentage of our daily workload.

Applying a systematic approach to software troubleshooting is the key to solving all problems. A systematic solution also works well in preventing problems in the first place.

Many of the common software problems that you will spend time solving can be prevented with proper preventive maintenance. Preventive maintenance tends to get neglected at many companies because technicians are too busy fixing problems. Spending some time on keeping those problems from occurring is a good investment of resources.

In this chapter, we'll look at applying the same troubleshooting methodology to common software problems. We'll also apply similar troubleshooting to security issues. First, we'll look at common symptoms of problems and their solutions. We'll then follow up with ways to deal with—and prevent—security-related issues.



Exam Objective 3.1, “Given a scenario, troubleshoot common Windows OS issues,” requires you to know the common troubleshooting of the Windows operating system.

Troubleshooting Common Microsoft Windows OS Problems

Windows is mind-bogglingly complex. Other operating systems are complex too, but the mere fact that Windows has nearly 60 million lines of code (and thousands of developers have worked on it!) makes you pause and shake your head. Fortunately, you just need to take a systematic approach to solving software-related issues.

Windows-based issues can be grouped into several categories based on their cause, such as boot problems, missing files (such as system files), configuration files, and virtual memory. If you're troubleshooting a boot problem, it's imperative that you understand the Windows boot process. Some common Windows problems don't fall into any category other than “common Windows problems.” We cover those in the following sections, followed by a discussion of the tools that can be used to fix them.

Symptoms/Issues

There are numerous “common symptoms/issues” that CompTIA asks you be familiar with for the exam. They range from the dreaded *Blue Screen of Death (BSOD)* to spontaneous restarts and everything in between. They are discussed in the following section.

Degraded Performance

The performance of your systems will inevitably slow down over time. This could be due to a multitude of causes, ranging from bad Windows Update patches to malware. Sluggish or slow performance is one of the hardest problems to solve on a Windows operating system, because many of the symptoms are related to each other.

The first step to solving the problem is identifying the component that is impacted by the performance issue. The following is a list of critical components that can be affected by slow performance:

CPU A symptom of poor CPU performance is the slow execution of applications. The operating system GUI will be unresponsive and sluggish. You may also hear your CPU fan running higher than normal if overheating starts to occur. CPU problems can be caused by an application that requires high CPU usage, such as movie rendering.

RAM A symptom of the operating system running out of RAM is high disk activity. As physical memory is used up, the less active pages of memory are swapped out to the disk drive. The symptoms of performance closely resemble CPU-related issues, where applications are slow in loading. RAM problems can be caused by too many applications being open at once or an application that has high RAM requirements, such as a database.

Disk A symptom of poor hard disk drive performance is the thrashing of the drive heads on the platters of the drive. Thrashing occurs when there is excessive movement of the drive arm to locate information on the drive. Disk problems can be caused by excessive fragmentation, high RAM usage, or a high volume of drive usage by applications, such as a video capture.


Network Symptoms of poor network performance are slow loading web pages, network applications that load slowly, and even timeouts. If you are using wireless, network issues can be caused by poor signal strength. If you are connected through Ethernet, poor network performance can be directly related to the local area network (LAN). The problems can also be outside of your network.

Graphics Symptoms of poor graphics (GPU) performance are usually related to slow-running video games and playback of videos. The frames per seconds (FPS) will be excessively low as the computer tries to render the screen. It is not common to have graphics-related issues unless you just downloaded the latest shoot-em-up game. Graphics-related problems are usually the hardest to solve because they require third-party tools by the graphics card vendor.

As you can see from the list of possibly affected components, many of the symptoms are closely related, such as RAM, CPU, and disk. The excessive usage of RAM can create performance symptoms with your hard disk drive. If left for a long period of time, these can both lead to an increase in CPU activity.

There are several tools that you can use to identify the problem area so that you can focus your attention on narrowing down the problem. The first tool you should start up is the Task Manager, as shown in Figure 7.1. You can launch Task Manager in several different ways, such as right-clicking the Start menu and selecting Task Manager, right-clicking the Taskbar and selecting Task Manager, or (my personal favorite) pressing Ctrl+Shift+Esc. The Performance tab will show you four of the five critical areas (detailed previously) on the left side. In this example, you can see that the processor is spiked out at almost 100 percent and all other systems are within tolerance.

FIGURE 7.1 The Performance tab in Task Manager



Now that you've isolated the problem to the critical area of CPU, you can narrow it down further by looking at the Processes tab, as shown in Figure 7.2. You can see that Google Chrome is using nearly 70 percent of the CPU. The other 30 percent is most likely distributed among other processes. By clicking the core area headings of CPU, Memory, Disk, and Network, you can sort usage from high to low or from low to high. In this

particular instance, the operating system was caught booting up, so that particular process was displaying high CPU.

FIGURE 7.2 The Processes tab in Task Manager

Name	Status	CPU	Memory	Disk	Network
Apps (9)					
Amcrest Surveillance Pro Appl...		4.5%	370.9 MB	0 MB/s	6.8 Mbps
Google Chrome (76)		69.6%	21,205.0 MB	0.1 MB/s	0 Mbps
Microsoft Excel (2)		0%	67.9 MB	0 MB/s	0 Mbps
Microsoft Word (2)		0.1%	163.7 MB	0 MB/s	0.1 Mbps
Photos (3)		0%	153.1 MB	0 MB/s	0 Mbps
SnippingTool.exe		0%	62.0 MB	0 MB/s	0 Mbps
Task Manager		0%	83.1 MB	0 MB/s	0 Mbps
VMware Workstation (32 bit) (5)		0%	3.3 MB	0 MB/s	0 Mbps
Windows Explorer		0%	210.1 MB	0 MB/s	0 Mbps
Background processes (83)					
AAC DRAM HAL (32 bit)		0%	0.1 MB	0 MB/s	0 Mbps
AAC MB HAL (32 bit)		0%	1.9 MB	0 MB/s	0 Mbps
AAC MB HAL (32 bit)		0%	0.1 MB	0 MB/s	0 Mbps
AacAmbientLighting.exe		0%	0.2 MB	0 MB/s	0 Mbps
AacKingstonDramHal_x64.exe		0%	0.1 MB	0 MB/s	0 Mbps
AacKingstonDramHal_x86.exe...		0%	0.1 MB	0 MB/s	0 Mbps
AcPowerNotification (32 bit)		0%	3.1 MB	0 MB/s	0 Mbps
Adobe Genuine Software Mon...		0%	0.1 MB	0 MB/s	0 Mbps
Antimalware Core Service		0%	3.8 MB	0 MB/s	0 Mbps
Antimalware Service Executable		0%	111.0 MB	0 MB/s	0 Mbps
Application Frame Host		0%	0.9 MB	0 MB/s	0 Mbps
ARMOURY CRATE Service		0%	4.4 MB	0 MB/s	0 Mbps
ARMOURY CRATE User Sessio...		0%	9.5 MB	0.1 MB/s	0 Mbps
ArmourySocketServer		0%	0.4 MB	0 MB/s	0 Mbps



Starting in Windows 8, the Task Manager gained a few tabs. One of those is called Details. The Details tab allows you to see the details of the processes, such as the user who executed the process, the process ID (PID), and the name of the process executable. You can sort by any of the headings by clicking them, as shown in Figure 7.3, just like you can on the Processes tab. You can even right-click the headings and add columns, such as peak memory usage and the exact command line, just to name a few.

Using *Resource Monitor*, you can get a much more detailed view than what is displayed in Task Manager. You can open Resource Monitor by clicking Start > All > Windows Tools > Resource Monitor, or you can click the Start menu and type **Resource Monitor**, and select it when it populates in the search results. Resource Monitor is shown in Figure 7.4. This tool allows you to read real-time performance data on every process on the operating system. Resource Monitor also allows you to sort details, the same as Task Manager. You can click each critical area and drill down to the performance issue.

FIGURE 7.3 The Details tab in Task Manager






FIGURE 7.4 Resource Monitor



A unique feature of Resource Monitor is the visualization of data. When you select a process on the upper view, Resource Monitor automatically filters the activity of the critical area, as shown in Figure 7.5. As you can see in this example, the Chrome browser processes have been selected and then the Network tab can be chosen to display the network activity and connections. The result is the isolation of network activity for this process. This can be done for any of the critical areas.

FIGURE 7.5 Selective isolation in Resource Monitor



Now that you've isolated the problem to an action or process in the operating system, you need to do the following:

1. Formulate a theory of probable cause.
2. Test the theory to determine the cause.
3. Establish a plan of action to resolve the problem and implement the solution.
4. Verify the full-system functionality and, if applicable, implement preventive measures.
5. Document findings, actions, and outcomes.

The theory of probable cause may be that a hardware upgrade is required due to a new version of the application that demands more resources. Or it could be as simple as the job the application is running is higher than normal in load. Remember to question the obvious and do the simple things, such as rebooting, to see if the problem goes away. It is often joked about that problems always go away after a reboot. It's not too far from the truth.

Sometimes a process is hung up and is affecting another process. A reboot sometimes fixes them both and the symptoms go away. More likely than not, the problem will still be there. This is when you need to start testing your theory of probable cause to determine the cause.

You might find that after running a large query the hard drive is extremely stressed. Your action plan might be to upgrade the hard drive or move the workload to a faster machine. In either case, you need to verify that the process is functioning the way the user expects it to.

If you determine that it's a certain report in the database, remind the user that the report takes time and maybe they should schedule it when the computer is not immediately needed. Or schedule the system for an upgrade of hardware to prevent these problems in the future.

Ultimately, you should document your finding so that other technicians do not waste time with the same issue. The more intricate the problem, the more time is wasted when you forget you've solved it already or don't remember the answer. Always document the actions taken, such as upgrades or changes to the process for the user. You should also note the outcome—whether it was successful or showed no immediate performance increase—so that if another technician is working on the same or a similar issue, they can gauge whether the solution is effective.

Boot Problems

With the introduction of Windows Vista, the boot process had changed from prior operating systems of Windows XP, 2000, and NT 4.0. We've used this current boot process introduced with Windows Vista all the way to today with Windows 11. The current boot process allows for the adoption of UEFI firmware.

In order to troubleshoot a failure to boot, you need to understand the complete boot process, starting with either the BIOS or the UEFI. The process is slightly different depending on which firmware you have on the motherboard. However, the outcome is the same: the hardware hands control over to the operating system so that the operating system can boot.

BIOS Legacy BIOS systems perform a *power-on self-test (POST)*, and then the BIOS bootstrap routine looks at the master boot record (MBR) at the beginning of the disk. The MBR then reads the boot sector on the first primary partition found. This boot sector then instructs the Boot Manager to load.

UEFI UEFI firmware will perform a similar POST. Then the UEFI bootstrap begins by loading drivers for the hardware. One of the differences is that UEFI can contain drivers that allow it to boot across a network or other nonstandard devices. Just like the legacy BIOS, the UEFI firmware looks at the MBR in the GUID Partition Table (GPT). The GPT defines a globally unique identifier (GUID) that points to a partition containing the Boot Manager. Therefore, UEFI firmware requires a partitioning scheme of GPT and cannot use the standard MBR partitioning scheme.

The initial boot sequence from hardware control to software control is almost identical in both BIOS and UEFI firmware. UEFI firmware does give you many more options, because

UEFI drivers can be loaded before control is handed over to the software. This allows UEFI to treat all locations containing an operating system the same. Up to the point at which the hardware hands control over to the software, there is no difference between a network boot and a hardware boot.



Windows 11 requires a UEFI firmware and therefore the boot process will follow the UEFI process. However, legacy MBR booting is still used with operating systems that you put on a USB drive, such as the Windows installation process. UEFI BIOS typically supports legacy BIOS booting, such as physical media, ISOs, or USB devices.

After control is handed over to the software, several files are used to complete the operating system bootup. The most important files are as follows:

Windows Boot Manager The Windows Boot Manager (BOOTMGR) bootstraps the system. In other words, this file starts the loading of an operating system on the computer.

BCD The Boot Configuration Data (BCD) holds information about operating systems installed on the computer, such as the location of the operating system files.

winload.efi winload.efi is the program used to boot Windows. It loads the operating system kernel (ntoskrnl.exe).

winresume.exe If the system is not starting fresh but resuming a previous session, then winresume.exe is called by BOOTMGR.

ntoskrnl.exe The Windows OS kernel is the heart of the operating system. The kernel is responsible for allowing applications shared access to the hardware through drivers.

ntbtlog.txt The Windows boot log stores a log of boot-time events. It is not enabled by default.

System Files In addition to the previously listed files, Windows needs a number of files from its system folders (e.g., SYSTEM and SYSTEM32), such as the hardware abstraction layer (hal.dll), the Session Manager (smss.exe), the user session (winlogon.exe), and the security subsystem (lsass.exe).

Numerous other dynamic link library (DLL) files are also required, but usually the lack of them or corruption of one of these files produces a noncritical error, whereas the absence of hal.dll causes the system to be nonfunctional.

We'll now look at the complete Windows boot process. It's a long and complicated process, but keep in mind that these are complex operating systems, providing you with a lot more functionality than older versions of Windows:

1. **Power On:** The system self-checks and enumerates hardware resources. Each machine has a different startup routine, called the POST (power-on self-test), which is executed

by the commands written to the motherboard of the computer. Newer PnP boards not only check memory and processors, but also poll the systems for other devices and peripherals.

2. **UEFI Firmware Initialization:** After the POST, the UEFI firmware continues to initialize the rest of the hardware components. This involves setting up the CPU, initializing the memory, and detecting connected storage devices. The UEFI firmware also prepares the system environment for the operating system to take over.
3. **Boot Manager:** Once the hardware initialization is complete, the UEFI firmware hands control over to the Windows Boot Manager, called `BOOTMGR`. The Boot Manager is responsible for managing the boot process and determining which operating system to load, especially if you have multiple operating systems installed.
4. **Boot Configuration Data (BCD):** `BOOTMGR` reads the boot configuration data (BCD) to get a list of boot options for the next step. The BCD contains multi-boot information or options on how the boot process should continue.
5. **Windows Boot Loader:** The Boot Manager then loads the Windows Boot Loader (`winload.efi`) from the EFI system partition. The Boot Loader is responsible for loading the Windows kernel and essential drivers needed to start the operating system.
6. **Suspended State:** If Windows is returning from a hibernated (suspended) state, `winresume.exe` is responsible for reading the `hiberfil.sys` file into memory and passes control to the kernel after this file is loaded.
7. **Kernel Initialization:** The Windows Boot Loader initializes the Windows kernel (`ntoskrnl.exe`), which is the core of the operating system. The kernel is responsible for managing system resources, such as memory, processes, and hardware interactions. This switches the system from real mode (which lacks multitasking, memory protection, and those things that make Windows so great) to protected mode (which offers memory protection, multitasking, and so on) and enables paging. Protected mode enables the system to address all the available physical memory.
8. **Hardware Abstraction Layer (HAL):** The Windows kernel loads the executive subsystems. *Executive subsystems* are software components that are required for the operating system to communicate with the hardware. One of the subsystems is the HAL, which provides a consistent interface for the operating system to interact with the hardware, ensuring compatibility across different hardware platforms.
9. **Session Manager:** The kernel loads the Session Manager. The Session Manager (`smss.exe`) is then started, which initializes system sessions and loads the Registry. The `HKEY_LOCAL_MACHINE\SYSTEM` Registry hive and device drivers are loaded. The drivers that load at this time serve as boot drivers, using an initial value called a *start value*.
10. **Winlogon:** `Winlogon.exe` loads. At this point, you are presented with the login screen. After you enter a username and password, you're taken to the Windows desktop.


11. **Explorer.exe:** Finally, the Windows Explorer (`explorer.exe`) process starts, providing the graphical user interface (GUI) for Windows 11. This includes the desktop, taskbar, and Start menu, allowing you to interact with the operating system and launch applications.

Now that you understand the boot process, let's look at how you can collect information to identify the problem. We'll consider this in two parts: hardware and software. The hardware process begins with the POST, and the software portion of the bootstrap begins with the BOOTMGR.


You can collect information from the BIOS/UEFI firmware boot with third-party system event log (SEL) viewers. However, it is very unlikely that you have a failure to boot because of a BIOS/UEFI firmware issue. It's not impossible, but it is highly unlikely.

To collect information on the software portion of the boot process loads, you can use boot logging. The `ntbtlog.txt` file is located at the base of the `C:\Windows` folder, as shown in Figure 7.6. Boot logging is off by default and needs to be turned on. To enable boot logging, issue the command `bcdedit /set {current} bootlog Yes`. You can also use the System Configuration utility (`msconfig.exe`) by selecting the Boot Log option on the Boot tab, as shown in Figure 7.7. Because the BCD is read by BOOTMGR, this point of the boot process is where logging would begin and the first entries would be the loading of the kernel.

FIGURE 7.6 The `ntbtlog.txt` file



```
Microsoft (R) Windows (R) Version 10.0 (Build 22621)
12 21 2024 12:33:49.500
BOOTLOG_LOADED \SystemRoot\system32\ntoskrnl.exe
BOOTLOG_LOADED \SystemRoot\system32\hal.dll
BOOTLOG_LOADED \SystemRoot\system32\kd.dll
BOOTLOG_LOADED \SystemRoot\system32\mcupdate_AuthenticAMD.dll
BOOTLOG_LOADED \SystemRoot\System32\drivers\CLFS.SYS
BOOTLOG_LOADED \SystemRoot\System32\drivers\tm.sys
BOOTLOG_LOADED \SystemRoot\system32\PSCHED.dll
BOOTLOG_LOADED \SystemRoot\system32\BOOTVID.dll
BOOTLOG_LOADED \SystemRoot\System32\drivers\FLTMRG.SYS
BOOTLOG_LOADED \SystemRoot\System32\drivers\msrpc.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\ksecdd.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\clipsp.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\cmimcext.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\werkernel.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\ntosexts.sys
BOOTLOG_LOADED \SystemRoot\system32\CI.dll
BOOTLOG_LOADED \SystemRoot\System32\drivers\cng.sys
BOOTLOG_LOADED \SystemRoot\system32\drivers\Wdf01000.sys
BOOTLOG_LOADED \SystemRoot\system32\drivers\WppRecorder.sys
BOOTLOG_LOADED \SystemRoot\system32\drivers\WDFLDR.SYS
BOOTLOG_LOADED \SystemRoot\System32\DriverStore\FileRepository\prm.inf_amd64_de435dc5c75d64a5\PRM.sys
BOOTLOG_LOADED \SystemRoot\System32\Drivers\acpiex.sys
BOOTLOG_LOADED \SystemRoot\system32\drivers\msseccore.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\ACPI.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\WMILIB.SYS
BOOTLOG_LOADED \SystemRoot\System32\drivers\msisadrv.sys
BOOTLOG_LOADED \SystemRoot\System32\drivers\pci.sys
```

FIGURE 7.7 System Configuration options for boot logging


Chances are, if you're having trouble booting into Windows, you won't be able to access the command prompt to issue `bcdedit` commands, nor will you be able to access `msconfig .exe`. Not to worry. You can still access logging by allowing the operating system to fail two times in a row. The third time, the computer will boot into the recovery console, as shown in Figure 7.8. From there, click **Troubleshooting > Advanced Options > Startup Settings > Restart**. When the computer restarts, it will boot into the Startup Settings menu, as shown in Figure 7.9. Of course, you won't be able to boot the computer to retrieve the files, but you can use the command prompt in the Windows Recovery Environment.

The idea is to collect information to identify the problem and, above all, to fix the problem. Sometimes you need to let Windows repair itself. The Windows Recovery Environment (WinRE) contains a Startup Repair option. Using the Startup Repair option is similar to issuing `bootrec /rebuildbcd` at the command prompt, which will rebuild the BCD.

If that fails, the ultimate solution might be to use the **Reset This PC** option in the Windows Recovery Environment or to install the operating system from scratch.

No OS Found

When it's reported that an operating system is missing, or "no OS is found," the first thing to check is that no media is in the machine (USB, DVD, CD, and so on). The system may be reading this media during boot before accessing the hard drive. If that is the case, remove the media and reboot. You should also change the BIOS/UEFI settings to boot from the hard drive before any other media to prevent this issue in the future.

FIGURE 7.8 Recovery Environment menu

If the problem is not as simple as removing the non-bootable media, then you may have to boot into the Windows Recovery Environment. This may be a challenge, because if the BIOS/UEFI cannot boot to the Windows Boot Manager, then the Windows Recovery Environment cannot be executed. The Boot Manager is responsible for executing the Windows Recovery Environment. You will have possibly two options to fix this. The first option is to use the vendor's recovery console. This option is dependent on the vendor supplying a recovery console that can be accessed via the BIOS/UEFI; not every vendor supplies this tool. The second option is to boot the installation media and choose Repair My PC when it first boots, as shown in Figure 7.10. Choosing this option will launch the Windows Recovery Environment booted from the installation media. You can then choose to repair the operating system by selecting Troubleshoot > Startup Repair. The Windows Recovery Environment will then attempt to repair the operating system.

System Instability

When an *application crashes*, you want to isolate the cause of the crash and solve it. The cause could be a compatibility issue, a hardware issue, or a host of other problems. One step to take early on is to look for updates/patches/fixes to the application released by the vendor. Be sure to try these updates on a test machine before rolling them out to all machines, and verify that they address the problem and do not introduce new problems.

FIGURE 7.9 Startup Settings menu






FIGURE 7.10 Windows 11 Repair My PC menu




One tool that is extremely helpful in identifying software problems is Reliability Monitor, shown in Figure 7.11. Reliability Monitor allows you to see application crashes and the times and dates they occurred. It also allows you to see which updates were installed before and after the crashes. You can use Reliability Monitor to narrow down whether other software is causing the issues and what led up to the crashes. To open Reliability Monitor click Start > All > Windows Tools > Control Panel > Security And Maintenance, open Maintenance, and click View Reliability History. Alternately, you can click the Start menu and type View Reliability History, then select it from the active search results.

FIGURE 7.11 Windows Reliability Monitor



In addition to Reliability Monitor, you can access the Windows event logs in *Event Viewer* for information about Microsoft-based application problems, as shown in Figure 7.12. All third-party vendors should log errors to the Windows event logs, but generally you will only find Microsoft products using these logs. In either case, you might find more information about why an application is crashing by looking at the Application log. You can open the Windows Event Viewer by clicking Start > All > Windows Tools > Event Viewer, or you can click the Start menu and type **Event Viewer**, then select it from the active search results.

FIGURE 7.12 Windows event logs in Event Viewer

The applications on the operating system are not the only elements captured by Reliability Monitor. Reliability Monitor also captures the overall stability of the operating system. It will allow you to see every reboot caused by a problem and even blue screens. The overall stability of the operating system is drawn as a graphical line inside the Reliability Monitor application. This allows you to look back and trace when a problem started.

In Exercise 7.1 you will review the Reliability Monitor built into Windows.

EXERCISE 7.1

Reviewing Reliability Monitor

This exercise is best performed on an operating system that has been running for some time. Although the operating system should be flawless and operate at a reliable level, you should see some fluctuation in reliability. If you don't, then try switching the view from days to weeks in the upper left of the app, above the graph.

1. Click the Start menu, type **View Reliability History**, and then click the Control Panel app result.
2. Make note of your current reliability score by looking at the graph on the rightmost section of the graph and comparing the location with the legend on the left side (1 through 10).


3. Make note of if and why your score has gone down over the days shown.
4. If there is a drop in reliability, click Day, and the lower section of the app will display the reliability details.
5. View the details and make note of the action.

As you click through the reliability events, make note of applications that constantly crash or that are problematic. These are the applications you should investigate further. Some of the solutions in the next section will lend some suggestion on how to fix them.

Blue Screen of Death

The Blue Screen of Death (BSOD)—not a technical term, by the way—is another way of describing the blue-screen error condition that occurs when Windows fails to boot properly or quits unexpectedly, as shown in Figure 7.13. If this happens during a boot, it is at this stage that the device drivers for the various pieces of hardware are installed/loaded. If your Windows GUI fails to start properly, more likely than not the problem is related to a misconfigured driver or misconfigured hardware.

FIGURE 7.13 Blue Screen of Death



You can try a few things if you believe that a driver is causing the problem. One is to try booting Windows into Safe Mode, which you can access via the Startup Settings in the Windows Recovery Environment. In Safe Mode, Windows loads only basic drivers, such as

a standard VGA video driver and the keyboard and mouse. After you've booted into Safe Mode, you can uninstall the driver that you think is causing the problem.

Another option is to boot into the Windows Recovery Environment and use System Restore, which will revert the system drivers back to the state they were in when the restore point was created. Bear in mind that a System Restore will not affect personal files, but it will remove applications, updates, and drivers.

The Windows Recovery Environment will automatically launch if there are two failed attempts to boot the operating system in 2 minutes. You can manually launch the Windows Recovery Environment by holding the Shift key as you reboot the operating system. Another way is to boot the installation media and select Repair My PC.

Frequent Shutdowns


There are many reasons why your computer might randomly turn off or shut down without warning. The problem is almost always related to faulty hardware or a faulty driver, but sometimes it can be as simple as tweaking your advanced power settings.

The first place to check is Event Viewer on the System tab. You should start combing through the logs and looking for the source of Kernel-Boot or Kernel-General. The entries will help you identify if the operating system was shut down properly or if it just suddenly lost power. Any time the operating system is shut down, the kernel will log an entry, and when it is powered back up, it will also log an entry. In addition to the Kernel-Boot and Kernel-General sources, you should also investigate the source of EventLog. These entries will be created if the EventLog service detects a dirty shutdown, such as when the power is removed. By checking the Event Viewer logs, you identify whether the problem is a hardware problem or whether the operating system is actually shutting itself down.

If you identify that the problem is a hardware issue, then the first step to resolving the problem is updating drivers. You should remove any autodetected drivers and reinstall the vendor's driver for the specific hardware. If the problem persists, then swapping known good hardware might help narrow down the issue.

If you determine that the operating system is shutting itself down, then the power settings should be checked, as shown in Figure 7.14. You can access the power settings by navigating to Start > Settings > System > Power > Screen and Sleep Timeouts. Alternately, you can access the Power Control Panel applet by opening Control Panel and selecting Power Options, then selecting Change Plan Settings on the selected power plan. Make sure that the PC is not going into sleep mode when plugged in. From this screen you can access the advanced power settings by clicking Change Plan Settings, then Change Advanced Power Settings. You will need to selectively tweak some of the settings and test your adjustments, such as turning off hard disk sleep, general sleep settings, and processor power management, to name a few. Tweak the timer or adjustment and then monitor the change from the setting to determine the fix.

FIGURE 7.14 Power Options



Services Not Starting

A service's failure to start is directly related either to another application installed with conflicting resources or to a misconfiguration of the service. In either case, the first place to start is the Event Viewer, as shown in Figure 7.15. The System log will display an Event ID of 7000 from the source of the Service Control Manager. The reason for the failure will vary, depending on the problem.

If a service is conflicting with another resource, we recommend that you reinstall the software that installed the service that is failing. Although this might break the conflicting application, it is probably the quickest way to find a conflicting resource.

If the service fails to start because of a misconfiguration, the most likely cause is the user account the service is configured to start with. If a misconfigured user account is the problem, you will see an Event ID of 7000 in Event Viewer, and the description will read that the service failed due to a login failure. You can verify the user configured to start the service in the Services properties, as shown in Figure 7.16. You open the properties of the service by right-clicking the Start menu, selecting Computer Management, then Services, right-clicking the service, selecting Properties, and finally selecting the Log On tab.

FIGURE 7.15 Service Control Manager events





FIGURE 7.16 Services properties



Make sure that the password for the user account has not changed and that the user account is not locked out. You can manually reset the password for the user and reenter the password in the Services properties. Also make sure that the account has the Log On As A Service right.



Real World Scenario

Did You Reboot Your Computer?

Quick quiz: You just got an error in Windows, and it appears that you are on the verge of a crash (of your application or the whole system). What do you do?

The first thing is to write down any error messages that appear. Then save your work (if possible) and reboot your computer.

Anyone who has called tech support, or who has been a tech support person, knows how demeaning the question, “Did you restart your computer?” can seem. Most people respond with an indignant, “Of course!” In reality, they might or might not have actually done it.

Whenever there’s a software problem, always, always reboot the computer before trying to troubleshoot. Often, the problem will disappear, and you’ll have just saved yourself half an hour of frustration. If the same problem reappears, then you know that you have work to do.

Why does rebooting help? When an application is running, it creates one or more temporary files that it uses to store information. It also stores information in memory (RAM). If a temporary file or information in RAM becomes corrupted (such as by application A writing its information into application B’s memory space), the original application can have problems. Rebooting will clear the memory registers and most often remove problematic temporary files, thus eliminating the issue.

It might sound trite, but the first axiom in troubleshooting software really is to reboot. Even if the user says they did, ask them to reboot again. (Tell them you want to see the opening screen for any possible error messages, or make up another good excuse.) If the problem doesn’t come back, it’s not a problem. If it does, then you can use your software skills to fix it.


Low Memory Warnings

Random access memory (RAM) is the computer’s physical memory. The more RAM you physically have installed in a computer, the more items you can have in the running foreground of the operating system. If you run out of physical memory, then processes that are

backgrounded (minimized) will be loaded into the *page file*, or *paging file*, on the hard drive. This is called paging, and it is totally normal to have a certain amount of paging happen during normal activity. The page file is actually hard drive space into which idle pieces of programs are placed, while other active parts of programs are kept in or swapped into main memory. The programs running in Windows believe that their information is still in RAM, but Windows has moved the data into near-line storage on the solid-state drive (SSD) or hard drive. When the application needs the information again, it is swapped back into RAM so that the processor can use it.

When system processes are at risk of not having enough memory free, you will see a warning message similar to the one shown in Figure 7.17. When this happens it means one of two things; the first is that you simply don't have enough physical RAM in the computer. The second is that a process is using a large amount of RAM that it normally doesn't need. The operating system is letting you know that it can't swap out any more pages of memory to the page file (virtual memory).

FIGURE 7.17 Low memory warning



The larger the page file, the fewer times the machine has swapped out the contents of what it is holding in memory. The maximum possible size of your page file depends on the amount of disk space that you have available on the drive where the page file is placed. Windows configures the minimum and maximum page file size automatically. If you want Windows to handle the size of the page file dynamically, you have to change the default setting by selecting System Managed Size in the Virtual Memory dialog box. We'll show you how to get there in a moment.


In Windows, the page file is called `pagefile.sys` and is located in the root directory of the drive on which you installed the OS files. The page file is a hidden file; to see the file in Windows File Explorer, you must have the Folder options configured to show hidden files. Typically, there's no reason to view the page file in the filesystem, because you'll use Control Panel to configure it. However, you may want to check its size, and in that case, you'd use Windows File Explorer.



The moral of the story: As with most things virtual, a page file is not nearly as good as actual RAM, but it is better than nothing.

To modify the default virtual memory settings, follow these steps: Click Start menu > Settings > About > Advanced System. In the Performance area, click Settings. Next, select the Advanced tab (yes, another Advanced tab), and then, in the Virtual Memory area, click Change and the Virtual Memory dialog box will open, as shown in Figure 7.18. Note that in addition to changing the page file's size and how Windows handles it, you can specify the drive on which you want to place the file.

FIGURE 7.18 Windows Virtual Memory settings




You should place the page file on a drive with plenty of empty space. As a general rule, try to keep 20 percent of your drive space free for the overhead of various elements of the OS, like the page file. Do not set the page file to an extremely small size. If you make the page file too small, the system can become unbootable—or at least unstable. In general, the page file should be at least 1.5x the amount of RAM in the machine.

USB Controller Resource Warnings

The *USB controller* is a hardware component on the motherboard that supplies both power and a data path for the devices connected. It is possible to plug in too many devices and overload the power the port can handle. Most USB 2.0 ports can handle five concurrent loads of 100 mA each, for a total of 500 mA. USB 3.0 can handle six concurrent loads of 150 mA each, for a total of 900 mA. If a device connected draws more than the allotted power, it will malfunction or irradicably disconnect.

The USB controller is also responsible for allotting the number of endpoints for the purpose of accepting data. The endpoints are equivalent to the number of lanes on a highway for cars (data) to travel at any given moment. If you plug in too many devices you can request more ports than are allotted for the USB controller and you will get an error, similar to Figure 7.19.

FIGURE 7.19 USB controller error




The easiest way to fix this issue is to move some USB devices around on the USB ports. You should move any devices that don't need USB 3.0 to USB 2.0 ports, such as keyboards and mice. Then ensure that your devices that require speed support are connected to USB 3.0 ports. If you are using a USB hub, make sure that you have connected it to a similar port. For example, if you are using a USB 2.0 hub, ensure it is connected to a USB 2.0 port and not a USB 3.0 port. More complex solutions are upgrading the driver to the latest driver the vendor supplies for the USB controller or just simply upgrading the hardware to a newer chipset.

Newer systems will come with a charging downstream port (CDP). These ports can provide higher current (up to 1.5 amps) to charge devices faster while still allowing data transfer. To put things in perspective, USB-C charging bricks can provide up to 100 watts (20 volts at 5 amps); these devices are also considered USB Power Delivery (PD). It is not common to find these types of ports on a laptop or computer, it is usually the interface that powers the device from the wall outlet.

Slow-Loading Profiles

A local profile is a group of settings for the user as well as their personal files. Local profiles can be slow to load because of items set to start when the profile is loaded. You can use Task Manager to selectively disable startup items, as shown in Figure 7.20. By the process of elimination and after several logouts and logins, you can narrow down the performance problem caused by slow-loading local profiles.

FIGURE 7.20 Startup items in Task Manager



Typically, local profiles will not slow down login tremendously. Roaming profiles that are located on a server typically cause slow loading of the profile. The local profiles don't need to traverse a network during login. Roaming profiles, on the other hand, need to traverse the network during login (load from the server) and logout (write back to the server).

There are some things you can do to alleviate the stress on the network and speed up the load time of network profiles. For example, you can save space by deleting temporary Internet files in both Edge and Internet Explorer. You can also save a tremendous amount of space—sometimes gigabytes—by deleting downloaded files. In addition to space traversing the network, login scripts, Group Policy processing, and services starting upon login can also contribute to slow-loading profiles.

Time Drift


The *real-time clock (RTC)* on the motherboard is responsible for maintaining the correct time. The RTC can drift over time and the computer can become faster or slower. When an operating system is running on a hypervisor, the problem is increased tenfold, since the RTC is actually emulated by the hypervisor. When the *time drifts* too far, you can have authentication problems. If the time drifts too far, certificates can also be invalidated and you'll have problems with web browsers.

Fortunately, the Windows operating system has addressed the problems of time drift by periodically querying a Network Time Protocol (NTP) server. You will need to ensure that

the client has the ability to contact the time server of `time.windows.com`, or you will need to configure a time server the client can reach.

You can verify that the NTP server is reachable by opening the Date & Time Settings screen and trying to update time. This can be performed by clicking the Start menu > Settings > Time & Language > Date & Time. Be sure that the Set Time Automatically option is selected. Then click Sync Now, under the Additional Settings heading. The operating system will attempt to call out to the NTP server and the results will be displayed in the dialog box, as shown in Figure 7.21.

FIGURE 7.21 Configuration of date & time



On virtual machines running on a hypervisor, the VMware guest tools or Hyper-V Integration Services will be responsible for emulating the RTC. These additions will also sync the VM's clock to the hypervisor host's time; this way, you only have to worry about the host's synchronization with an NTP server.

Troubleshooting Steps

Now that we've covered common symptoms of Microsoft Windows OS problems, let's look at solutions that you can implement to solve those problems.

Rebooting

Rebooting a system often takes care of problems, for a multitude of reasons. One of the top reasons is that it allows the operating system to terminate hung processes gracefully. After the operating system reboots, the processes are normally restarted. An added bonus is that only the applications the user requires are relaunched. The goal is to fix the user's problem with minimal disruption. A reboot can't be done every time, especially if the user could lose work as a result.

Depending on the circumstances, when reproducing a problem, one of the first things to do is reboot. It serves an important purpose of isolating the problem so that you can reproduce it. For example, if you've isolated the problem to Excel not scrolling properly when a web browser is open, you should reboot and try to replicate the problem. If you discover the reboot fixed the problem, then you've solved the problem. However, if the problem still exists, you've now isolated the problem further by eliminating other programs that could have been hung in the background affecting this problem. The steps of rebooting and then opening Excel can also be used to verify when you've solved the problem.

Restarting Services

Services normally don't need to be restarted. On occasion, however, a change is made that requires that a service be restarted to reflect the change. Services should be restarted if they crash, of course. Although this is rare, it still happens from time to time. If a service crashes, you can restart the service in the Computer Management MMC by selecting Services, then right-clicking Service and choosing Start, as shown in Figure 7.22. You can use the same method to restart a running service. The Computer Management MMC can be started by right-clicking the Start menu and selecting Computer Management.

Services can be configured to automatically start in the event of failure on the Recovery tab of the Services properties, as shown in Figure 7.23. For example, by default; the Print Spooler service is set to restart on the first and second failure, but after that it will remain stopped. The reset counter can be set for a number of days, and the service can be started after a specific number of minutes after its failure. You can even have the computer restart or run a program in the event a service fails.

Fixing Application Issues

If an application is crashing and acting erratically, it may be due to another application that has overwritten critical files used by the application, or the files may have become corrupted. In either case, choosing to repair an application will validate that it is installed properly and the process will replace any missing critical files for the application. Data files and configuration files will not be touched while the application is being repaired; only critical files (such as DLLs) will be checked and repaired.

FIGURE 7.22 Manually starting a service






FIGURE 7.23 Service recovery



You can repair an application by right-clicking the Start menu, selecting Installed Apps, finding the application, clicking the three dots to the right of the application, selecting

Advanced Options, then selecting Repair under the Reset heading, as shown in Figure 7.24. The application's installer will launch and start to repair the application. Older applications might not have the Advanced Options menu; in that case, after clicking the three dots, select Modify. If the installer supports a repair, you will be presented with the option to repair the installation or remove it.

FIGURE 7.24 Repairing an application



If a repair does not fix the application, then you should perform a complete uninstall and reinstallation of the application. When you uninstall the application, the uninstaller should remove configuration files that could be causing the issue. Some applications require you to manually remove the configuration files, which are often in a folder of your profile, such as AppData. You should always contact the vendor for the specific

folders to remove to ensure the application has been completely uninstalled. Once you've completely uninstalled the application, you can reinstall the application and test to see if the problem is fixed.

If after a repair or uninstall and reinstall, the application is still broken, it could be caused by a bug. You can update the application to get the latest fixes. However, it is always recommended to identify the problem and then cross-reference it with the vendor's change log for the application.



Using the attitude “I can’t break it any more than it is broke” might work for most applications. However, just blindly updating an application may cause other problems and complicate the original issue. Always try to verify if the problem is a known bug and update the application accordingly.

Application Requirements and Resources

Applications require a certain amount of RAM, storage space, and CPU speed. Some applications may also require an SSD hard drive, or a GPU with a specific speed and amount of VRAM. When searching for a solution for a problem, you should first verify the requirements for the application based on the vendor's requirements. This will establish an expectation of performance for the application on the given hardware. If the requirement is higher than the given hardware, then you will need to scale the hardware up by adding resources.

Although the hardware might meet the requirements of the application, most often the application is not the only application running on the hardware. This must be taken into consideration when trying to solve a problem.

The performance of the application might also need to be addressed. Often vendors will publish the minimum specification for the application requirements. However, when you speak with support their recommendations for your workload might be very different. This is another aspect of the solution that must be taken into consideration.

Computers are built with a finite amount of resources, such as RAM, CPU, and storage. After reviewing the application requirements and weighing the considerations of application coexistence with other applications and the current load of the application in question, you may decide to add more resources. Adding resources such as RAM, CPU, and storage is considered to be scaling up the hardware. Before adding resources, you should document the performance and the utilization of the resources. Then after you've added the resources, you should compare the current performance and utilization of resources by the application.



The System File Checker (sfc.exe) was covered in Volume 2 Chapter 2, “Windows Configuration” and maps to Exam Objective 1.5, “Given a scenario, use the appropriate Microsoft command-line tools.”

System File Checks

When a feature of the operating system stops functioning or errs in a manner that makes you suspect corruption, the System File Checker tool can scan and replace critical files. The System File Checker is launched from the command line with the command `sfc.exe` and performs a myriad of functions. For example, you can execute the command `sfc.exe /verifyonly`, and the System File Checker will inspect all the critical files and verify integrity only. You can also supply the command `sfc.exe /scannow` and the tool will scan and repair any files that fail the integrity check. You can perform the same task on individual files such as `kernel32.dll`. The System File Checker tool will also allow offline repair and checks.

Using the System File Checker tool on the operating system to be inspected is pretty straightforward, as described previously. However, if the operating system can't boot, you can use the command in offline mode as well. You simply boot the Windows Recovery Environment locally from installation media or some other source. Once the Windows Recovery Environment is booted, select Advanced Options and open a Command Prompt as Administrator. Then enter `sfc.exe /scannow /offbootdir=C:\ /offwindir=C:\Windows`, assuming that Windows is installed on the C: drive.


Resetting Windows

As previously discussed, the System File Checker utility will verify the integrity and replace any corrupted files. However, the utility will only replace files if they fail an integrity check. An alternate method to ensure the operating system is properly installed is to perform a repair installation of Windows. The repair installation will reinstall all files from source media regardless of their integrity. The repair installation will leave all applications and user files in place.

To initiate a repair installation of Windows, click Start > System > Recovery > Reset This PC. You can then select to Keep My Files or Remove Everything. If you select Keep My Files, then the operating system will reinstall itself not touching any of your data. The latter of the two will completely reset the PC and all data. The next option you have is to perform the reset from a Cloud Download or Local Reinstall. The Cloud Download option will download and reinstall Windows from the cloud. The Local Reinstall option will be used to locate files to reinstall the operating system. You will then be prompted to confirm your settings and clicking Reset will begin the reset (reinstallation) of the operating system, as shown in Figure 7.25.

Restoring Windows

Almost everyone, no matter how hard they've tried to keep their computer running properly, will experience a computer crash at some point. Many of the ways to get your computer back up and running (such as reinstalling the operating system) take a lot of time. In Windows, System Restore allows you to create restore points to make recovery of the operating system easier.

FIGURE 7.25 Preparing to reset

A *restore point* is a copy of your system configuration at a given point in time. Restore points are created a couple different ways:


- You can manually create them yourself (which is highly recommended before you make any significant changes to the system, such as installing new drivers).
- During the installation of some programs, a restore point is created before the installation; that way, if the installation fails, you can “roll back” the system to a preinstallation configuration.

Restore points are useful for when Windows fails to boot but the computer appears to be fine otherwise, or if Windows doesn’t seem to be acting right and you think it was because of a recent configuration change.

It’s important to note that in Windows 11, the automatic system recovery option is disabled by default. It must be turned on manually. We will cover how to do so in Exercise 7.2.

Microsoft is moving toward unifying all settings under the Settings app for Windows 11. This differs from prior operating systems and the legacy Control Panel app. System Restore is one of those settings that can be opened only from the legacy Control Panel app. To open System Restore, click Start, type **Control Panel** and select it from the results, and then click Recovery. This will open to a list of Control Panel choices. Select Configure System Restore. This will open the System Properties, and you can then select Configure and select Turn On System Protection, as shown in Figure 7.26.

FIGURE 7.26 System Protection





If you need to use a restore point and Windows won't boot, the system will boot in the *Windows Recovery Environment*, in which a full system restore can be performed.

The other option is to select how much disk space is available for System Restore. The less disk space you make available, the fewer restore points you will be able to retain. If you have multiple hard drives, you can allocate a different amount of space to each drive.

Exercise 7.2 demonstrates how to create a restore point manually in Windows.

EXERCISE 7.2

Manually Creating a Restore Point in Windows

1. Click the Start menu, type **Control Panel**, and then click the Control Panel app result.
2. Click Recovery, if the menu item does not appear, adjust your View by from Category to Large Icons.
3. Click Configure System Restore. The System Properties dialog box will open.
4. Select Local Disk (C:) (System), and then click Configure.
5. Click the Turn On System Protection radio button, and then click OK.
6. Click Create below the Configure button.
7. Type a name for the restore point.
8. When the process is finished, click OK, and then exit out of the Control Panel windows.

Reimaging/Reinstalling the Operating System

In certain situations, a problem may require you to reinstall software. The time required to uninstall and reinstall the software can sometimes exceed the time it takes to reimage the operating system with the software preinstalled. Reimaging the computer will depend on whether you use operating system images or load each computer by hand.

If your organization does not use a standardized image for its computers, you can use the Windows Recovery Environment and select the Reset This PC option. If the computers have a preinstalled image, you can use the System Image Recovery option to reload the operating system. You can select this option by holding down the Shift key as you reboot the operating system, then choosing Advanced Options after the reboot, then selecting System Image Recovery. Depending on the type of computer you have, it may have a proprietary process for recovering system images.



A factory recovery partition may exist on a machine. A vendor has a great deal of freedom with what it puts here, but usually it contains copies of drivers and preinstalled programs. This partition—if it exists—may serve as a last resort for stabilizing a system if all else has failed.

Windows Updates

Occasionally, applying an update will fix a problem, mainly because that is what updates do: they fix problems. Once you've identified that applying an update is the solution, you need to download, distribute, and install the update. Luckily, by default Windows 10/11 automatically installs updates for the operating system to keep you up to date and problem-free.


In large-scale networks, the organization may employ a corporate patch-management solution. Microsoft offers a free patch-management solution called Windows Server Update Services (WSUS). Microsoft also sells a licensed solution called Microsoft Endpoint Configuration Manager (MECM), which performs many other functions in addition to patch management. If an update is required and your organization uses one of these products, the patch must be approved, downloaded, and deployed. Third-party patch-management solutions may also be used in your organization. Third-party solutions are usually specific to an application or suite of applications, such as Adobe or Autodesk.

In small office, home office (SOHO) environments and small network environments, the update may be a one-off installation for a specific application. In this case, the update just needs to be downloaded and installed, per the vendor instructions. Always make sure to have a plan to roll back from a bad update. Turning on System Protection before the update is a good idea. If an update fails, you can simply use System Restore to restore the operating system to a prior point in time.

Very rarely you will find that a Microsoft or third-party update has created a problem on the operating system. When this happens, it's pretty easy to roll back updates by uninstalling them. Simply open the Settings app, select Windows Update, then Update History, then Uninstall Updates, and finally select the update and choose Uninstall, as shown in Figure 7.27.



An alternative to uninstalling an update is to use System Restore to revert to an earlier time before the update. However, System Protection must have been turned on before the update was installed. System Protection in Windows 11 is turned off by default. You can turn it on by clicking the Start menu, then typing **Recovery** and selecting it, then clicking Configure System Restore, selecting the System drive, clicking Configure, selecting Turn On System Protection, and finally clicking OK.


FIGURE 7.27 Uninstalling an update

Rolling Back Device Drivers

When you isolate a hardware problem to a faulty device driver, it is sometimes necessary to roll back the current driver to a prior version. This action will roll back the driver to the original version detected by Windows, also called the *out-of-box driver*. In some cases, it may roll back to a generic driver, which reduces functionality until a proper driver is installed.

This process can be completed with these steps:

1. Right-click the Start menu.
2. Select Device Manager.
3. Select the device.
4. Right-click and select Properties from the context menu.
5. Select the Driver tab.
6. Click Roll Back Driver.
7. Provide the reason for rolling back the driver.
8. Click Yes, as shown in Figure 7.28.

FIGURE 7.28 Rolling back a driver

When the rollback is complete, you should reboot the computer before testing to see if it fixed the issue.

Rebuilding Windows Profiles

When a problem has been determined to be a profile-related issue, it is necessary to reset the Windows profile. An example of a profile-related issue is a slow loading profile. This is often related to the accumulation of non-critical files that bog down the login process. A profile reset allows the user to start with a clean profile without all the bloat of these files. When performing this action, ensure that the user's data is backed up. It is best to keep an entire copy of the profile before resetting it. The following are the most common places data is kept by the operating system:

- Contacts
- Desktop Items
- Documents
- Favorites
- Links

- Music
- Pictures
- Videos
- Saved Games (optionally)
- Downloads (optionally)
- 3D Objects (optionally)

To back up a local profile, log into an administrative account (other than the one you are backing up), and then copy the profile under `C:\Users` to a new location. Do not move the profile, because the operating system references it in the Registry.

You can then reset a local profile on the Advanced tab of System Properties, as shown in Figure 7.29. You can access the User Profiles dialog box by following these steps:


1. Click the Start menu.
2. Type **System**.
3. Click Advanced System Settings in the search results.
4. Select the Advanced tab.
5. Select Settings (User Profiles).
6. Select the profile you want to reset.
7. Click Delete, and answer Yes to confirm the deletion.
8. Log in as the user. Windows will create a new profile.

The user's files can then be manually copied over. The Profile dialog box also allows you to view the overall size of a user's local or remote profile, so it also helps in the troubleshooting process.

You can also use that procedure to delete a roaming user profile that has been left on the Windows operating system. However, performing the procedure on a roaming profile will not reset the profile. You will only remove the profile to clear space. To reset a network-based roaming profile, perform the following steps:

1. Ensure the user is logged out completely.
2. Delete all local copies of the user's profile left on any machine.
3. Navigate to the network location containing the user's profile and rename the folder.
4. Log the user into the machine on which you have deleted the locally cached copy.
5. Copy any useful items back to the user's profile.
6. Log the user out to ensure the roaming profile is saved back to the user's profile.

FIGURE 7.29 Deleting a local user profile



User profile management can become complicated, depending on how it is implemented. Exam Objective 3.1, “Given a scenario, troubleshoot common Windows OS issues,” requires you to have a basic understanding of how to remedy a slow profile load. For additional information on user profiles outside of the objective, visit:

[https://learn.microsoft.com/en-us/previous-versions/windows/desktop/legacy/bb776892\(v=vs.85\)](https://learn.microsoft.com/en-us/previous-versions/windows/desktop/legacy/bb776892(v=vs.85))



Exam Objective 3.4, “Given a scenario, troubleshoot common personal computer (PC) security issues,” requires you to have a basic understanding of common operating system symptoms and browser-related symptoms as they relate to security.

Troubleshooting Security Issues

There are a number of topics CompTIA expects you to know for the 220–1202 exam as it pertains to security issues. Many of these issues also appear in other CompTIA certification exams, such as the CompTIA Security+ and other exams that have a security component. Rest assured that for the 220–1202 exam, you do not need to know the depth of content as if you were preparing for the CompTIA Security+ exam. However, you should be familiar with the following symptoms when determining if a security-related problem has occurred:

- You are unable to access the network.
- Desktop alerts mysteriously appear.
- You receive false alerts regarding antivirus protection.
- Altered system or personal files are discovered.
- Missing or renamed files are discovered.
- Unwanted notifications appear within the operating system.
- Operating systems updates fail.
- You receive random or frequent pop-ups.
- You receive certificate warnings.
- Your browser is redirected to websites.

This list is by no means exhaustive. What is an absolute is the fact that you should immediately rectify the security issue or quarantine the system if you experience even one of these symptoms. In the following sections we will cover the aforementioned security-related symptoms, as well as explore some causes for these symptoms.



It cannot be overstated that establishing security policies and procedures, updating your operating systems, updating your applications, and updating your network devices are all good measures to help eliminate potential security problems.

Common Symptoms

The operating system is by far the largest attack vector for threat agents. The software installed on the operating system and the files that the operating system stores are the perfect mixture of targets. Threat agents can target an unpatched application that contains a vulnerability. A threat agent can also sneak a file that is infected with malware onto the operating system. In the following section we will identify symptoms of common security issues along with their possible causes.

The solutions for the symptoms vary significantly. However, the overall goal is to keep your operating system protected. Protecting the operating system can be achieved by keeping

the OS current on all patches. The applications that are installed on the operating system should be kept current with patches as well. Antimalware software should be installed, and best practices should be employed, as previously discussed in past chapters.

Network Connectivity Issues

If your computer is hooked up to a network, you need to know when your computer is not functioning properly on the network and what to do about it. In most cases, the problem can be attributed to either a malfunctioning network interface card (NIC) or improperly installed network software. The biggest indicator in Windows that some component of the network software is nonfunctional is that you can't log in to the network or access any network service. To fix this problem, you must first fix the underlying hardware problem (if one exists) and then properly install or configure the network software.

In some situations, network connectivity issues can be related to security threats. Although you might not seem to have network connectivity, the NIC might be working fine, while the real problem is a malicious program that has crashed or that it is not operating as the creator of the malicious program intends. The malware will act as a *proxy* for the network traffic. This type of malware is usually intent on stealing credentials or banking information. However, it can also be used to inject ads and cause *browser redirection*.

Not all malware that causes network connectivity issues acts as a proxy. Some malware changes network settings, such as your DNS servers. This type of malware will cause browser redirections by controlling what you resolve through its DNS. It is also common for malware to change your system proxy so that all requests go through the threat agent's remote proxy.

False Alerts and Hoaxes

Users have plenty of real viruses and other issues to worry about, yet some people find it entertaining to issue phony threats disguised as security alerts to keep people on their toes. Some of the more popular *hoaxes* that have been passed around are the Goodtimes and the Irina viruses. Millions of users received emails about these two viruses, and the symptoms sounded awful. The mention of these two hoaxes serves to outline the most well-known hoaxes. Since these came out, there have been many different hoaxes, most of which were not as well known as these.

Both of these warnings claimed that the viruses would do things that are impossible to accomplish with a virus. When you receive a virus warning, you can verify its authenticity by looking on the website of the antivirus software you use, or you can go to several public systems. One of the most helpful sites to visit to get the status of the latest threats is the website for the *Cybersecurity & Infrastructure Secure Agency (CISA)*. CISA monitors and tracks threats and provides regular reports on this site.


When you receive an email that you suspect is a hoax, check the CISA site before forwarding the message to anyone else. The creator of the hoax wants to create widespread panic, and if you blindly forward the message to coworkers and acquaintances, you're helping the creator accomplish this task. For example, any email that includes "forward to all your friends" is a candidate for research. Disregarding the hoax allows it to die a

quick death and keeps users focused on productive tasks. Any concept that spreads quickly through the Internet is referred to as a *meme*.

Desktop Alerts

A desktop alert is a notification or dialog box that is crafted to look like it was generated by the operating system. This is a crafty way of social engineering the user into becoming a victim. The malware is crafted to generate a pop-up box that states there is a security error detected and that you should call Microsoft or Windows Support right away, as shown in Figure 7.30. When you call the number given, you are calling scammers who will try to sell you software you don't need.

FIGURE 7.30 Malware-generated call-in alert




Social engineering is not the only method of a threat agent. The threat agent can generate realistic operating system dialog boxes that coax you into downloading and installing malware, such as shown in Figure 7.31. The average person might just figure it's time to update their software—it even states that, by downloading the software (malware), you agree to the EULA.

OS Notifications

Another really popular method of distributing malware is by using browser push notification messages. The user will browse to a malicious site and then the user will be coaxed into allowing push notifications for the site. Once this is allowed, the site can push notifications to the operating system and spawn a notification that looks like it's coming from the

operating system. Use of the operating system notifications is a well-known attack aimed at coaxing the user into installing malware or pushing advertising to the user.

FIGURE 7.31 Malware-generated download alert



In some cases, the initial installation of the malware is prompted by a browser push notification. After the user installs the malware, it might start prompting deals of the day or other ads. This type of malware is considered adware, and it is becoming rare compared to other types of malware covered in this chapter.

Outside of user education, antimalware software can be used to prevent this type of threat. However, user education is much more effective. A routine review of sites allowed to send notifications should be performed periodically. A routine review of installed applications should be performed as well to ensure that malware has not been installed at some point.

Rogue Antivirus

One clever way of spreading a virus is to disguise it so that it looks like an antivirus program. When it alerts the user to a fictitious problem, the user then begins interacting with the program and allowing the rogue program to do all sorts of damage. One of the trickiest things for threat agents to do is to make the program look as if it came from a trusted source—such as Microsoft—and mimic the Windows Notification Center interface enough to fool an unsuspecting user. The notification might show that a new download or update is waiting for you to install it. It may even notify you that your antivirus software is disabled and needs attention.

Education is the only way to combat a rogue antivirus. You should arm yourself with the knowledge of current antivirus programs. You can achieve this knowledge by reading consumer articles on the latest and greatest antivirus and antimalware applications. You should also pass that education to others in your family and organization. This can easily be achieved by detailing which antivirus and antimalware you have installed. Then there will be no confusion when a notification or pop-up occurs stating it's time to install the rogue anti-virus malware.



Threat agents are designing malware to look very realistic and mimic big-name software. If you are unsure of the source and do not trust it, visit the site and redownload the software. Do your homework, as the threat agents are doing their homework.

Renamed System Files/Disappearing Files/Permission Changes/Access Denied

Threat actors that create malware have a number of methods by which they can wreak havoc on a system. One of the simplest ways is to delete key system files and replace them with malicious copies. When this occurs, the user can no longer perform the operation associated with the file, such as printing, saving, and so on. When malware is embedded on an operating system and gains privilege level access, it is known as a *root kit*. Once the operating system is infected with the malware, the threat actors will comb through files looking for sensitive information that they can ransom.

Just as harmful as deleting files, the malware can rename the files or change the permissions associated with them. This could prevent the user from accessing the files or even copying them off to an uninfected system. When an operating system is infected with *ransomware* the malware will encrypt the files. The files might also disappear from the user's normal view and the ransom request may be placed in the parent directory. The mode of operation for most ransomware is to rename the files with a unique extension as the malware encrypts the files. The ransom note is then placed in every folder so the user has instructions on payment for decrypting the files.

Starting with Windows Vista, Microsoft enabled the User Account Control (UAC) by default. This change to the operating system greatly reduced the number of attempts to use elevated privileges and definitely made it more difficult to change system files. In addition to enabling the UAC, Microsoft removed the Modify NTFS permission from system files for the Administrator account. Only the *Trusted Installer* (Windows Update) has access to modify these files; even the System (operating system) permissions are Read and Execute. If that wasn't enough, a self-healing service watches for files changed and replaces them with trusted versions. The System File Checker (SFC) is a user tool that can be used to manually heal missing or modified system files. Malware can maliciously modify files and, in some cases, cause them to go missing. The System File Checker was covered in Volume 2 Chapter 3, "Windows Administration," and was also discussed earlier in this chapter.

OS Update Failures

Failed updates for Windows—assuming they aren't caused by connectivity issues—can often be traced to misconfigured settings. These settings can also cause the operating system to report that an update needs to be installed when it has already been installed. The best solution is to find the error code being reported in Windows Update Troubleshooter, solve the problem, and download the update.

Recent versions of Windows 11 include a troubleshooting utility. To access the utility, click the Start menu and select the Settings app. Once the Setting app opens, select System, then click Troubleshoot, then Other Troubleshooters, and finally choose Windows Update.

Microsoft has also published some common problems that should be checked. You can access the document at <https://support.microsoft.com/en-us/windows/troubleshoot-problems-updating-windows-188c2b0f-10a7-d72f-65b8-32d177eb136c>. The solutions Microsoft recommends in the documents are as follows:

- Reboot.
- Verify free space.
- Run Windows Update again.
- Update third-party drivers.
- Disconnect external hardware.
- Check Device Manager for errors.
- Perform a system restore or repair installation.

Browser-Related Symptoms

The web browser is the most used application on the operating system. It's so popular that Google has made an operating system around their Chrome web browser called the ChromeOS. On Windows, the web browser is just an application like any other application on the operating system. However, it is the easiest way a threat agent can access your operating system. Therefore, in this section we will cover some of the most common symptoms you may observe related to the web browser.

Pop-Ups

Pop-ups (also commonly known as popups) are both frustrating and chancy. When a user visits a website and another instance (either another tab or another browser window) opens in the foreground, it is called a pop-up; if it opens in the background, it is called a *pop-under*. Both pop-ups and pop-under are pages or sites that you did not specifically request and that may only display ads or bring up applets that should be avoided.

Most modern web browsers come standard with a pop-up blocker and by default they block all pop-ups. If you visit a link that contains a pop-up, the browser will notify you that it has blocked it. If the pop-up is on a trusted website, you will have the option to allow pop-ups for the site.

Threat agents have found other creative ways to pop up ads or referred content. Through the use of JavaScript they can serve overlays over the original web page. In older web browsers, the JavaScript could even hold you hostage at the malicious page. However, most newer browsers limit the access of JavaScript, and you can always close the web page.


If you continually receive pop-ups or overlays, then you may be infected with malware or a rogue page is minimized serving the pop-ups/overlays. A reboot should clear the problem, but it is also best to scan your operating system with antivirus/antimalware software.

Certificate Issues

There are several problems that plague *digital certificates*. Of the two major problems, one is related to the proper setting of time and date and the other is trust related. The time on the host should always be checked along with the expiration of the SSL certificate. If the certificate is expired, this will cause problems.

On the other hand, when an untrusted SSL certificate is encountered, the web browser will alert you that the SSL certificate is not valid, as shown in Figure 7.32. Every web browser comes with a list of trusted certificate publishers. If a certificate is issued to a website or is not trusted, a warning box will come up preventing you from visiting the site. You can click through the warning prompt and visit the site anyway, but the address bar will still read “Not secure” or display an unlocked lock icon during your visit.


FIGURE 7.32 An untrusted SSL certificate warning



The problem should always be investigated further, since information entered in the site could be intercepted if the site was hacked. The first step to diagnose is checking the hostname in the URL. All certificates must match the hostname in the URL that they are issued for. If you tried accessing the site by the IP address, this warning is benign and can

be disregarded. However, if you entered the correct hostname, then the certificate should be inspected. Every web browser is different, but every web browser will let you view the certificate. In Figure 7.33 you can see that the certificate has been self-signed.

FIGURE 7.33 A self-signed certificate



Both the Issued To and Issued By fields in the certificate are the same. This is common when the website is in development, but it is not normal once the website has been placed into production. It is also common on network management equipment that allows configuration through a web page. Often the management web page will use a self-signed certificate. For this purpose, the certificate can be imported into your trusted publisher certificate store so that it can be trusted in the future.

Browser Redirection

Pharming is a form of redirection in which traffic intended for one host is sent to another. This can be accomplished on a small scale by changing entries in the hosts file and on a large scale by changing entries in a DNS server, also known as *DNS poisoning*. In either case, when a user attempts to go to a site, they are redirected to another site. For example, suppose Illegitimate Company ABC creates a site to look exactly like the site for Giant Bank XYZ. The pharming is done (using either redirect method) and users trying to reach Giant

Bank XYZ are tricked into going to Illegitimate Company ABC's site, which looks enough like what they are used to seeing that they give their username and password.

As soon as Giant Bank XYZ realizes that the traffic is being redirected, it will immediately move to stop it. But while Illegitimate Company ABC will be shut down, it was able to collect data for the length of time that the redirection occurred, which could vary from minutes to days.

Another form of browser redirection is called affiliate redirection. This type of browser redirection can be very subtle. For example, when you search for a product and click the link in the results, the malware will redirect your browser to the intended site with an affiliate link attached. Now anything you purchase will credit a commission to the person who redirected the browser with the affiliate link. This malware is usually related to an unscrupulous plug-in in the browser.

Because an attacker can use many different tactics to launch browser redirection, the mitigation is not straightforward. However, implementing end-user education, maintaining updates for browsers and operating systems, and ensuring that your antimalware/antivirus software is up to date are best practices to protect against browser redirection.

Degraded Browser Performance

Degraded browser performance is a frustrating experience. The symptoms are often slow loading times, frequent crashes, and unresponsive tabs. This can be caused by a number of factors, including an excessive number of open tabs, an outdated browser, or the browser needs a restart to complete the upgrade process. However, there are more insidious causes of degraded performance—for example, the presence of malicious software, such as malware, cryptominers, and web pages running malicious JavaScript, and malicious plug-ins or add-ons.

Malicious proxies are a common cause of degraded browser performance. These proxies intercept and redirect your Internet traffic through a server controlled by a threat actor. This can lead to slower browsing speeds, as your data takes a longer route to reach its destination. Additionally, malicious proxies can be used to inject unwanted ads, steal sensitive information, or even install further malware on your system. You should regularly monitor your proxy settings and ensure they haven't been altered without your knowledge.

Malware can significantly impact browser performance. Cryptominers use your computer's resources to mine cryptocurrencies without your consent, leading to high CPU usage, increased power consumption, and overall system slowdowns. This not only affects your browsing experience but can also shorten the lifespan of your hardware, running it 100% all the time without your knowledge. Malware also comes in the form of adware, which generates intrusive ads and pop-ups. This of course results in degrading performance and potentially exposes you to more malicious content.

Malicious web pages running JavaScript are another culprit behind degraded browser performance. These web pages are often in the form of pop-unders that open behind your main browser window. The malicious JavaScript web pages often run scripts that consume system resources. They can be used to track your browsing habits, display unwanted ads, or

even execute malicious code. Because they run in the background, you might not immediately notice their presence, but their impact on performance can be significant. You should use an ad blocker and antimalware software to mitigate these issues. It's also important to keep your browser and security software up to date, to combat these shady attempts by threat actors.

Malicious plug-ins are another common cause of degraded browser performance. These plug-ins are often disguised as useful tools or extensions. Once installed, they can hijack your browser, redirecting you to unwanted websites, displaying intrusive ads, or even stealing your personal information. Just like any browser-related malware, they can also consume significant system resources, leading to slowdowns and crashes. To protect yourself, install plug-ins only from trusted sources, regularly review and remove any unnecessary extensions, and use security software to detect and block malicious plug-ins. In addition, you should monitor the permissions plug-ins have. It's a red flag for a plug-in to have complete access to see all web activity or even access to your filesystem.



Exam Objective 2.6, "Given a scenario, implement procedures for basic small office/home office (SOHO) malware removal," requires you to know how to remove malware on a small network such as a small office or home office.

Best Practices for Malware Removal

Best practices for malware removal are a key objective for the 220-1202 exam. The best way to think about this is as a seven-item list of what CompTIA wants you to consider when approaching a possible malware infestation. The following discussion presents the information that you need to know.

1. Investigate and Verify Malware Symptoms


Before doing anything major, it is imperative first to be sure that you are dealing with the right issue. If you suspect malware, try to investigate the type (spyware, virus, and so on) and look for the proof needed to substantiate that it is indeed the culprit.

You first need to identify the problem. This can be done with a multitude of tools, but hopefully your antivirus/antimalware software will be the first tool that helps to identify the problem. If the antivirus/antimalware software fails to identify the problem, then other third-party tools must be used.

Earlier in this chapter, in the section "Troubleshooting Common Microsoft Windows OS Problems," we introduced you to Resource Monitor to isolate performance problems. A similar tool, called Process Explorer, can be downloaded from Microsoft Sysinternals. This tool allows a different visualization from what Resource Monitor provides, as

shown in Figure 7.34. You can see the process list on the operating system; in this case, there is a process called `regsvr32.exe`. When you look more closely, you can see that it is creating network traffic and is very active on the operating system. The process is actually a ransomware application calling out to command-and-control servers. It is sneakily disguising itself as the `regsvr32.exe` utility, which is normally used to register DLLs.

FIGURE 7.34 Process Explorer



Unfortunately, this lone example will not give you the expertise of a professional virus/malware hunter. However, it provides just one of many examples of third-party software that can help you detect and identify viruses and malware running on a computer.

Many built-in tools, such as `netstat.exe`, can also provide assistance. For example, the `netstat -nab` command enables you to view all the processes on the operating system and their network connections. Using the `netstat -nab` command is how it was identified that something looked wrong with the `regsvr32.exe` process; otherwise, the command would have looked like any other process on the operating system.

In addition to applications that can identify viruses and malware, third-party websites can aid in detection. One such website is VirusTotal (www.virustotal.com). VirusTotal allows users to upload potentially unsafe applications. Their service will scan the applications against more than 70 antivirus engines and report if the signature is found. It's a valuable tool to validate that you've found an application on your operating system that is malicious. Many tools, such as Process Explorer, can even check against the VirusTotal database.

2. Quarantine Infected Systems

Once you have confirmed that a virus or malware is at hand, then quarantine the infected system to prevent it from spreading the virus or malware to other systems. Bear in mind that the virus or malware can spread in any number of ways, including through a network connection, email, and so on. The quarantine needs to be complete enough to prevent any spread.


Ransomware is probably the biggest risk, since it will spread through a network rapidly and encrypt files in its path. The ransom is usually equivalent to the number of files or the total size of files. In either case, over the past eight years it has made headline news, as it has taken down extremely large companies. In one instance, the Petya ransomware even took down most of the computers in Ukraine, along with those of several other countries.

If an infected system is discovered and needs further analysis, it should be quarantined from the network and put into an isolated network. This hot network is a place where it can be studied further without repercussions to the operational network.

3. Disable System Restore in Windows

This is a necessary step because you do not want to have the infected system create a restore point—or return to one—where the infection exists. System Protection in Windows 10/11 is turned off by default. However, threat actors often turn System Protection on, so they can automatically restore your system back to when it was infected. It is best practice to disable System Protection, just in case the threat actor turned it on. You can disable System Protection with these steps:

1. Click the Start menu.
2. Type **Recovery** and select it from the results.
3. Choose Configure System Restore.
4. Select the system drive and click Configure.
5. Disable System Protection.
6. Click Delete (Disk Space Usage).
7. Click Continue (confirmation), then Close, OK, and finally Yes (confirmation), as shown in Figure 7.35.

FIGURE 7.35 System Protection


Most ransomware will dump your restore points for you, since you could potentially recover the operating system and then recover files using your previous versions.

4. Remediate Infected Systems

The steps taken here depend on the type of virus or malware with which you're dealing, but they should include updating antivirus and antimalware software with the latest definitions


and using the appropriate scan and removal techniques. You can update *Microsoft Defender* from the Windows Security screen; to access it, click Start > Settings > Privacy & Security > Windows Security > Open Windows Security > Virus & Threat Protection, and select Protection Updates under the Virus & Threat Protections Updates heading. You can also click the blue shield in the task tray (Notification area) in the lower-right corner of the desktop, which will open the Windows Security screen. You can then click Virus & Threat Protection, then Protection Updates under the Virus & Threat Protections Updates heading. Then click Check For Updates, as shown in Figure 7.36.

FIGURE 7.36 Microsoft Defender Antivirus Security updates




Depending on the type of virus or malware, you may need to boot into Safe Mode or the Windows Recovery Environment (as discussed earlier in this chapter). However, the remediation of the virus or malware will be different for each situation. Microsoft Defender Security can automatically perform an offline scan. To perform an offline scan, inside of the Windows Security screen click Virus & Threat Protection, click Scan Options, and select Microsoft Defender Offline Scan, as shown in Figure 7.37.

After you confirm that you will save your work by clicking Scan in the confirmation dialog box, the UAC will prompt you to answer Yes, and then Windows will reboot. The Windows Recovery Environment will boot, and Windows Defender Antivirus will run, as shown in Figure 7.38.

FIGURE 7.37 Microsoft Defender Antivirus (offline scan)

In some situations, such as in a ransomware attack, no remediation can be performed because the user files are encrypted. In these cases, the malware should be removed from the operating system, and then the user data must be restored from a backup. The unfortunate and terrifying fact when it comes to ransomware is that there will be loss of work.

In many instances, remediating the virus or malware is impossible because no one knows for sure what the virus or malware actually does. Antivirus researchers can document the delivery system that a virus or malware uses to enter your system. You can then patch the vulnerability, which is part of the remediation process. What antivirus research cannot do most of the time is document the payload of a virus or malware. This is because most of the time the payload is encrypted and changed, depending on the need of its creator. In these cases, the remediation might be to sanitize the drive and reinstall the operating system from an image or manually install it.


FIGURE 7.38 An offline Microsoft Defender Antivirus scan

5. Schedule Scans and Run Updates

The odds of the system never being confronted by malware again are slim. To reduce the chances of it being infected again, schedule scans and updates to run regularly. Most antimalware programs can be configured to run automatically at specific intervals; however, should you encounter one that does not have such a feature, you can run it through Task Scheduler.

Microsoft Defender Security is scheduled to automatically scan the operating system during idle times. However, if you want to schedule a scan, you can use Task Scheduler:

1. Click the Start menu.
2. Type **Task Scheduler** and then select Task Scheduler from the results.
3. Open the Task Scheduler Library.
4. Select Microsoft, then Windows.
5. Select Windows Defender, and double-click Windows Defender Scheduled Scan.
6. Select the Triggers tab.
7. Click New on the Triggers tab, then select Weekly and choose the day of the week in the New Trigger dialog box, as shown in Figure 7.39.
8. Click OK, You'll need to click OK again after the New Trigger dialog box closes.

FIGURE 7.39 Creating a Windows Defender Security scheduled scan

Windows Defender Security is scheduled to automatically download updates during the Windows Update check, which is daily. If you require the latest updates, use either the Check For Updates option in the Windows Update settings or the Check For Updates option in the Microsoft Defender Security Center.

6. Enable System Restore and Create a Restore Point in Windows

Once everything is working properly, it is important to create restore points again, should a future problem occur and you need to revert back. You can enable System Protection by following these steps:

1. Click the Start menu.
2. Type **Restore** and select it from the results.
3. Click **Create A Restore Point**.
4. Select the System drive and click **Configure**.
5. Click **Turn On System Protection**, and click **OK**.

You can then manually create a restore point by clicking **Create** in the System Protection dialog box, typing a description (such as **after remediation - date**), clicking **Close** (in the confirmation dialog box), and clicking **OK** to close the System Properties.

7. Educate the End User

Education should always be viewed as the final step. The end user needs to understand what led to the malware infestation and what to avoid, or look for, in the future to keep it from happening again. This training can be formal training in a classroom setting, or it can be an online training in which the user must participate and answer questions.

It is common for large companies to require annual or biannual end-user training for threats. It is becoming more common for training to be done online, and a number of companies offer this as a service. It is not uncommon for a company to send a phishing attempt to their employees. When an employee falls for the phishing attempt, they are automatically signed up for mandatory training. Incentives are also common, such as the first employee who notifies the IT department of the phishing attempt gets a gift card.



Exam Objective 3.2, “Given a scenario, troubleshoot common mobile OS and application issues,” requires you to know how to troubleshoot mobile OS issues. These include issues of an application failing to launch, an application failing to close or crashes, an application failing to update, an application failing to install, is slow to respond, the OS failing to update, battery life issues, random reboots, connectivity issues, near-field communication (NFC), and the screen does not autorotate.

Troubleshooting Mobile OS Issues

As mobile devices have been rapidly replacing the desktop and laptop machines that used to rule the workplace, the equipment an administrator must maintain has now evolved to cover a plethora of options. This section focuses on common mobile OS and application issues and some of the tools that can be used to work with them. A subsequent section will look at the same topics with more of a focus on security.

In the following sections, we will cover many common symptoms of problems that are common with mobile OSs and applications. In this section, we cover how to identify application symptoms as they appear in the objectives.

Application Problems

Mobile devices are generally error free and function fine. However, as we install applications to the mobile OS, we introduce potential problems. This happens because a mobile application is generally developed for the nominal platform. The application is also usually only tested on one or two devices as the developer sees fit. The developer can’t account for every make and model of mobile device, and this is generally why we see application problems. In this section we will cover the most common application problems you may encounter with mobile devices. You will notice that they are all somewhat connected and usually have the same steps to rectify the problem.

Application Launch Issues

If an application (app) does not load (fails to launch), it could be attributed to a multitude of reasons. One common reason is that the application is still running in the background and is not really loading; it is just becoming a foreground application and becoming the application in focus. When you close an application, it sometimes doesn't close all the way down and free up memory. Instead, it gets moved to the background and is technically still running.

The first thing you should try if an application is not loading is to force-quit the application. To force-quit an application on an Android device, press the tab view (usually the leftmost soft button) and then swipe the application left or right to close. On Apple devices, double-tap the Home button or swipe up on the application you want to close.

Another common problem related to applications not loading is that sometimes the cache associated with the application is corrupted. This usually happens right after an application upgrades itself, which is all the time on mobile devices. On an Android phone, you can clear the application's cache by tapping Settings, tapping Apps, choosing the application, tapping Storage, and finally tapping Clear Cache. On Apple devices, tap Settings, tap General, tap iPhone Storage, choose the application, and finally tap Offload App and confirm it by tapping Offload App again. Clearing the cache will not affect the majority of the application's storage.

In many cases, an application will not allow you to clear its cache. If the option is not there, or after you've cleared the cache the application is still not loading, uninstall and reinstall the application. This option will (should) remove any data associated with the app.

You can remove Android applications by tapping Settings, tapping Apps, choosing the application, and tapping Uninstall. You can then visit the Google Play store and reinstall the application.

Apple is even simpler. All you need to do is tap and hold on an icon until all the icons dance back and forth. An X will be displayed in the upper-left corner of the application icon. Simply tap the X to uninstall the application. You can then visit the App Store and reinstall the application.

Application Fails to Close/Crashes

Another less common problem with applications is that they crash or fail to close out unexpectedly. This issue is less common since most applications run fine, but occasionally you will have a mobile application just close or crash out on you. This is frustrating, because it usually happens when you most need it, and it is generally an intermittent problem. This is how this particular problem differentiates itself from an app not loading; it doesn't crash every time.

The recommendations are the same for a crashing or closing application. Ultimately, you need to find the series of events or steps to trigger the bug and crash or close the application. Once you've reproduced the problem, it's time to try to fix the issue by doing one or more of the following:

- Force-quit the application.
- Clear the application cache.

- Clear the application data.
- Uninstall the application.
- Reboot the device.
- Reinstall the application.

If none of these solutions works, then it may be time to check the vendor's site for any similar problems (and solutions) encountered by others. Support for applications on mobile devices is normally forum or community support. However, some paid mobile applications have email-based support. Describe the issue, the device make and model, steps to reproduce the problem, and other applications installed on the device.

Application Fails to Update

You may have one or more applications on your mobile device that fail to update. Generally, the developers will backward support an old application until it doesn't make sense anymore, because of new features. The developer's expectation is that their application will be updated on all the mobile devices so that features still operate as expected. So the user will most likely see crashes, closes, or other erratic behavior if the application is not up to date.

The Google Play store and the Apple App Store manage the purchase, installation, and upgrade of applications. This upgrade maintenance on installed applications happens in the background without the user ever knowing it even happens. However, from time to time you may encounter an application that does not want to upgrade automatically. The first troubleshooting step should be to try to manually upgrade it from the Play Store or the App Store.

If manually updating the application does not work, then there are a few other steps you can take to troubleshoot the application, such as force-quitting the application and rebooting the phone to close any applications that may be stuck in memory. The next step is to temporarily disable any antivirus or antimalware software installed on the device. Then try to upgrade the application manually from the Play Store or the App Store.

Another consideration is to make sure that you are connected to the Internet via Wi-Fi. Most app store applications will treat cellular data as a metered connection and will not automatically update applications. Also make sure that the app store is configured to automatically update applications.

If all else fails and the application still doesn't want to upgrade, you can try to uninstall the application and reinstall it. You can follow the same guidelines for a crashing or closing application. Before uninstalling the application, make sure you check the compatibility for the latest version of the application. You could uninstall the application and find you can't reinstall it, because your device does not meet the minimum specifications. This could also be the original problem, which would explain why the application won't update.

Application Fails to Install

There are several reasons why a mobile app fails to install, and understanding them can help you troubleshoot the issue effectively. One of the most common reasons is insufficient storage space on your device. Mobile apps require a certain amount of free space to install.

The make and model of newer phones have a fixed amount of storage. If your device is running low on storage, the installation process may fail. To fix this problem, you can simply free up space by deleting unnecessary files, apps, or media. Additionally, moving some of your data to cloud storage can help create the needed space for the new application.

Another reason for app installation failure could be compatibility issues. Not all applications are designed to work on every device or operating system version. Typically, the mobile app ecosystem is responsible for weeding out (so-of-speak) apps that are not compatible with your mobile device. However, if the app shows up in the app store and the installation fails, it might not be compatible with your device's hardware or software. To resolve this, check the app's requirements in the app store and ensure your device meets them. Updating your device's operating system to the latest version can also help, as it may include necessary updates that improve compatibility with newer apps.

Network connectivity problems can also cause application installation failures. A stable Internet connection is required for downloading and installing apps, especially larger ones. If your connection is weak or intermittent, the download might be interrupted, leading to a failed installation. To address this, ensure you have a strong and stable Wi-Fi connection. If you're using mobile data, make sure you have sufficient data allowance and a good signal.

Performance Issues

There are a number of reasons you can have performance issues with a mobile device. Most performance issues are directly related to the applications on the device. For example, an application may use too much processing time and it could cause poor battery life and performance. A group of applications can use all available RAM and starve the unit for processing space. In the following we cover the various performance problems you may encounter with a mobile device.

Slow Performance

Slow performance is almost always related to RAM usage. Mobile operating systems operate the same as conventional desktop or laptop operating systems. The only real difference between traditional operating systems and mobile devices is the default action for an application is not to close it but to put it into the background. As programs are loaded into RAM, they allocate a percentage for their variables and inner workings. When RAM is filled up, the mobile device will swap background memory pages onto the built-in storage; this process is similar to the page file process. This slows down the device because now its focus is on clearing up memory for competing applications.

Fortunately, most mobile devices allow you to see the RAM usage at a glance and over a longer period of time. On an Android device, tap Settings > Battery And Device Care > Memory. You will see the memory usage for the device, along with each application and its own usage. You can also clear up memory from here, which basically just closes the applications. Unfortunately, you cannot monitor RAM usage on an Apple mobile device, but a simple soft reset works just as well.

It is uncommon to have a performance problem attributed to high CPU on a mobile device. That is not to say it doesn't happen; it's just uncommon. To narrow down problems with a particular application performing slowly, reboot the device and launch only that particular application to isolate and monitor its performance.

Frozen System

Frozen is a silly term that technicians use when something is not functioning or responding. It doesn't really describe the temperature, just the functionality—a block of ice.

If the system is frozen (not responding to a single thing), it will appear with the same symptoms as a nonresponsive touchscreen. One of the ways you can differentiate between a frozen system/lockup and a nonresponsive touchscreen is if the device will soft-reset. If the device will not soft-reset, then a hard reset might need to be performed. You should research the hard-reset procedure for the type of device you are trying to hard-reset. For example, Apple has several different types of procedures depending on the model and generation of the device. Samsung also has hard-reset procedures that differ based on model of device.

If the restart does not work, plug in the device, let it charge (an hour or more is recommended), and then try to restart it. The power level in the battery can sometimes be so low that performance is turned all the way down, and the device will appear unresponsive and frozen.

Random Reboots

Random reboots and restarts could be a symptom of a hardware issue. They can also be related to a problem in the operating system. When a problem is intermittent or random, it is very hard to diagnose. However, by checking the following you can at least rule out some of the most common culprits of random reboots:

Battery Health The most obvious culprit is the battery that powers the phone. Check to make sure that your phone has a good charge and the battery is not swollen. If your phone allows for the battery to be replaced, you should clean the contacts and replace the battery, then monitor for reboots.

Update The operating system should be on the current revision of software. If it isn't, then it should be updated to the latest. Stability issues are often addressed in updates for the operating system.

Storage Check the storage on the device and make sure that it is not over 90 percent. If it is over 90 percent, then you should clean up some space and monitor the device. You can do this by clearing data from applications, clearing application cache, and uninstalling unused applications.

Running Applications Check the applications running on the operating system. This should be checked 10 minutes after a reboot so that you can see the applications that are set to automatically start on boot. You should disable or uninstall any application

that is rarely used and automatically starts on boot.

Auto Restart On some Android phones there is an Auto Restart feature that should be turned off by default. This feature should be checked because it could have been erroneously turned on and might causing the random reboots. Apple mobile devices do not have this feature.

Factory Reset The dreaded factory reset is a last-ditch effort. It will, however, rule out the device's operating system from being the culprit of random reboots. You should also be selective in the software that you reinstall. A cyclical process of install and monitor, then install more and monitor, may help identify an application that is affecting the stability of the device.

OS Fails to Update

When an OS fails to update for a mobile device, there can be a number of reasons for the issue. However, as you will see, the common troubleshooting steps are no different than those for any other application on the mobile device.

Reboot A reboot of the mobile device is always recommended, since a process could be preventing the update from installing. A reboot of the operating system also allows for memory to be freed up since insufficient memory could be preventing the update.

Compatibility You should check to make sure that the operating system update is compatible with the mobile device hardware. It is common to find that a hardware device is only supported for 5 to 7 years, at which time it is no longer compatible with future operating system upgrades.

Storage Make sure that there is enough storage space on the mobile device to accommodate the update. When an update is to be applied to the device, at some point both operating systems will need to exist: the upgrade and the current operating system.

Connectivity Be sure that you have a Wi-Fi connection and that you are currently connected. Many operating systems will not download the update over a metered connection, such as cellular. A bad Internet connection can also prevent an OS update from completing.

If none of these suggestions allow you to successfully update the mobile device, then you can try to manually install the operating system update via the over-the-air (OTA) update. This will require a computer, and depending on the vendor, this option may not work. The vendor needs to support manual installation of updates with their third-party utility. Apple supports the iTunes application, which will work with the device to upgrade the operating system. For more information, visit <https://support.apple.com/en-us/108964>.

Extremely Short Battery Life

Batteries never last as long as you would like. Apple defines battery life as the amount of time a device runs before it needs to be recharged (as opposed to battery life span, which is the amount of time a battery lasts before it needs to be replaced). Tips for increasing battery life include keeping OS updates applied (they may include energy-saving patches), avoiding ambient temperatures that are too high or too low, letting the screen automatically dim, and turning off *location-based services*. You should also disconnect peripherals and quit applications not in use. Wi-Fi, for example, uses power when enabled, even if you are not using it to connect to the network. The following is a list of tips to conserve battery life:

- **Enable Power Saving Mode:** Most devices have a power-saving mode that reduces performance and limits background activity to extend battery life.
- **Turn Off Unused Features:** Disable Bluetooth, Wi-Fi, GPS, and NFC when not in use. These features consume a lot of power even when idle.
- **Use Airplane Mode:** When you don't need to be connected, turning on Airplane Mode can save a lot of battery by cutting off all radio signals.
- **Update Apps Regularly:** Developers often release updates to optimize battery usage. Make sure your apps are up to date.
- **Limit Background Processes:** Restrict background data usage for apps that don't need to be constantly updated.
- **Charge Smartly:** Try to keep your battery level between 40% and 80%. Avoid letting it drop to 0% or charging it to 100% frequently.
- **Use Dark Mode:** If your device has an OLED or AMOLED screen, using dark mode can save battery life as these screens use less power to display black pixels.
- **Close Unnecessary Apps:** Make sure to close apps that you are not using to prevent them from running in the background.

Outside of the preceding usage tips, sometimes battery life is attributed to a performance problem. High RAM usage can shorten battery life because power is expended on moving pages of memory in and out of RAM. This is done to keep the foreground applications running.

Overheating

When most mobile devices get too warm, they will tell you that they need to cool down before they can continue to be used, and they will automatically take measures to protect themselves (turning off features, closing apps, and so on).

One of the most concerning reasons is the lithium-ion (Li-ion) battery contained inside these devices. When a Li-ion battery gets too hot, you risk explosion or fire from a situation called *thermal runaway*. This is where the battery starts to get so hot that the separator inside the battery melts and causes a chain reaction.

Luckily, mobile devices shut down when they get too hot from ambient temperatures or internal temperature from the CPU. One of the best ways to prevent overheating is to avoid ambient temperatures that are too hot. Avoid having the device in direct sunlight for extended time periods, in a hot car on a summer day, or on top of a heat source. When the device does overheat, you can often help it cool down quicker by removing any protective case that may be there—and putting it back on later.

Mobile Device Overheating

When a mobile device overheats, be sure to understand why it is overheating. Was the ambient temperature too hot? Was the CPU usage too high and indicative of another problem? Was it a combination of the two, or did you just leave it in the direct hot sun?

When a lithium-ion battery is exposed to heat (even in small amounts), its useful life span will degrade. Most phones today have a built-in battery, which is not user-replaceable. When the battery goes bad, either you are tied to a charger every hour or it's time for a new phone.

Connectivity Issues

We rely heavily on our mobile devices, and the day-to-day functionality of our mobile devices relies heavily on connectivity to the outside world. The connectivity to the provider network (cellular) is generally supported by the provider, such as Verizon, T-Mobile, or AT&T, just to name a few. When you have a problem with the cellular network, make sure that you have coverage and a fresh reboot. However, outside of the obvious you are best to escalate the problem to your provider. The day-to-day connectivity we are responsible for as A+ technicians will be covered in the following section.

Intermittent Wireless

There are a number of reasons why intermittent wireless connections can occur, but the two most common are lack of a good signal and interference. Increasing the number of *wireless access points (WAPs)* for coverage, or being closer to them, can address the lack of a good signal.

Interference can be addressed by reducing the number of devices competing for the same channel. In many instances, however, the interference may be coming from an external source, such as a microwave oven or even a Bluetooth device on the 2.4 GHz band. To avoid common interference of this nature, use an SSID that is dedicated to the 5 GHz band. Using the 5 GHz band won't guarantee you an interference-free connection, since radar operates in this band. However, you will have better odds of selecting a channel without interference. In an effort to reduce interference and speed up wireless connectivity, 802.11ax/802.11be has been developed to use a 6 GHz band. The Wi-Fi Alliance officially released Wi-Fi 7 in 2024.

Another common problem with intermittent wireless is the auto-reconnection feature for the SSID. Your phone normally goes into a sleep mode every so often. This is normal and saves battery life. One of the first devices to sleep for battery conservation is the wireless circuitry. When you power on your phone, the wireless circuitry needs to associate with your WAP, which will happen unless you do not tell it to automatically reconnect. You can verify your wireless SSID and the auto-reconnection settings on Android by tapping Settings, tapping Connections, tapping Wi-Fi, tapping the current SSID, tapping Edit, and then making sure Auto Reconnect is selected. On an Apple device, tap Settings on your Home screen, tap Wi-Fi, tap the blue circled I next to your current SSID, and make sure that Auto-Join is on.



The auto-reconnect features can open you up to security concerns, such as the evil twin attack. The evil twin attack involves two access points; one is operated by an organization, and the other is operated by the threat agent. The threat agent will assign the same SSID to their AP, then they will send disassociation frames to the organization's AP and its clients. This forces the clients to disconnect, and if auto-reconnect is on they will reconnect to the evil twin operated by the threat agent. Information is then relayed to the network and actively sniffed for usernames and passwords.

No Wireless Connectivity

A common cause of a lack of wireless connectivity is that the wireless radio has been turned off. It happens from time to time, when an application that controls the Wi-Fi doesn't turn it back on. On an Android phone, swipe down from the status bar, then tap the wireless icon to make sure it is lit up. On Apple devices tap Settings, then tap Wi-Fi, and make sure that the slider is turned to the right and lit up in green.

Another common cause of a lack of wireless connectivity is for a device to be in *Airplane mode*. When a mobile device is in Airplane mode, all the radios for the cellular network of the provider, Wi-Fi, Bluetooth, and *near-field communication* (NFC) will be turned off. This function was created so that in one tap you could comply with the Federal Aviation Administration (FAA) or the European Aviation Safety Agency (EASA).




On newer iOS devices, such as iPhones, the Bluetooth is not turned off when you enter Airplane mode.

To make sure that your device is not in Airplane mode, look on the upper status bar that displays your cellular strength. If a plane appears there, then your phone is in Airplane mode. On Android, you can swipe down, then tap the icon of the airplane so that it's no longer lit up. On Apple devices, tap Settings, then Airplane Mode, and then tap the slider so that it is no longer lit up. On both Android and Apple devices, both the cellular service and wireless networks will be restored after Airplane mode is disabled.

Bluetooth Connectivity

The lack of Bluetooth connectivity can also be attributed to the use of Airplane mode, or Bluetooth can just be turned off. So, be sure to check this setting in addition to Airplane mode. On Android devices, you can swipe down, then tap the Bluetooth icon if it is not lit up. On Apple devices, go to Settings, then tap Bluetooth if it is not lit up. Depending on the phone and version of the operating system, the Bluetooth icons will be displayed on your upper status bar and will look similar to Figure 7.40.

FIGURE 7.40 Bluetooth status icons



Lack of Bluetooth connectivity can also be caused when a device is not turned on and/or has an improper setting for pairing. A common pairing issue is not having the proper Bluetooth passcode entered for the device. Each device, when paired, has a specific code from the vendor. Most vendors use a common code, such as 1234, but the code could also be 0000, or any combination, so it's best to check the vendor's documentation for Bluetooth pairing information.

To pair or re-pair a device, first ensure that the device is turned on and that it's discoverable. (Consult the vendor's documentation, as necessary.) On Android devices, tap Settings, then tap Connections, then tap Bluetooth (the phone will immediately start scanning for discoverable devices), select the available device, and enter the passcode from the vendor's documentation. On Apple devices, pairing can be performed by tapping Settings, tapping Bluetooth, tapping the device name, and entering the passcode from the vendor's documentation.

Near-Field Communication (NFC)

Near-field communication (NFC) is a short-distance wireless communication protocol. NFC is built into many mobile devices for the application of payment systems, such as Google Pay and Apple Pay. In addition, NFC is used for data exchange between mobile devices. A use case for this application is the transfer process when a new mobile device is purchased. You can simply tap the two devices together, and the new mobile device pulls the information from the existing mobile device. NFC nominally requires a distance of 4 centimeters or less to operate.

The first thing to check on the mobile device is that Airplane mode is not on. Airplane mode will impede the functionality of NFC, because NFC uses electromagnetic radio fields to enable communications between the phone and the NFC device.

The next thing to check is that it is not the reader. However, when attempting to pay for something and Google Pay or Apple Pay isn't working, anxiety often builds. Troubleshooting between you and the cashier is usually the last thing that comes to mind. However, briefly asking the cashier if anyone else has had an issue today can rule out the reader.

The case on the mobile device can also interfere with NFC communications. If the case is a ruggedized case and has an aluminum back, it could impede the NFC signal. Simply popping it out of the case can rule out the phone case as part of the problem.

Signing out of the mobile payment system will sometimes rectify the problem. In the process of signing back in, it will also validate if the provider's network is down, as this can often be a problem for mobile payment systems. Another troubleshooting step is selecting another credit card in the mobile payment system if you have multiple cards available in the app.




You can learn more about all of the various indicators for Apple devices at <https://support.apple.com/en-us/118203>. Android is an open operating system and customized by the phone vendor. Samsung has a similar article at www.samsung.com/uk/support/mobile-devices/what-do-the-different-indicator-icons-mean-on-your-galaxy-phone.

Autorotate Issues

The autorotate function allows a phone to switch between portrait mode and landscape mode by sensing how you are holding the phone. Autorotate is a feature of convenience, because no matter which way you are holding the phone, you can read the information displayed. This of course is assuming you have the screen facing you.

The first item to check is that you do not have autorotate turned off or locked. On the Android operating system there are several different ways to check this, depending on the vendor and the Android version, so it is best to check your specific model of phone. On Apple devices, this can be checked by swiping down from the top-right corner of your screen. When the Control Center opens, look for the circle with the lock in the middle and make sure it is not enabled. This is the Portrait Orientation Lock button, and if it is set, as shown in Figure 7.41, the screen will not autorotate.

If the autorotate function is not turned off or locked to a specific orientation, then you should suspect an application has possibly locked the orientation. A quick reboot will close out all running applications that could have a lock on the autorotate function. The reboot will also reset the autorotate service, in case it has crashed.

FIGURE 7.41 The Portrait Orientation Lock button

If closing the applications and rebooting the device does not remedy the issue, then you should suspect a hardware issue. There are third-party tools available that allow you to test the sensors. Ultimately, the service center can verify that a sensor is bad and malfunctioning.



Exam Objective 3.3, “Given a scenario, troubleshoot common mobile OS and application security issues,” requires you to know how to troubleshoot mobile OS security issues. These include issues of security concerns and common symptoms of security issues.

Troubleshooting Mobile Security Issues

The preceding section—and its corresponding objectives—looked at mobile devices and focused on common OS and application issues; this section builds on that and focuses on security-related issues. Once again, it looks at security concerns and common symptoms, differing only in that there is more of a focus on security. It needs to be pointed out, though, that CompTIA is stretching the definition of the word *security* to include more scenarios than many would typically consider. A fair number of the issues that appear in this section would have fit easily in the preceding section.

Security Concerns

As it pertains to mobile devices there are a number of security concerns that you should be aware of. These concerns are the same for personal devices as they are for organizationally owned devices. Understanding these concerns will help you secure mobile devices and allow you to be more knowledgeable about the consequences.

Application Unofficial Sources

Unofficial sources of applications for Apple devices, referred to as third-party app stores provide an alternative to the official Apple App Store. Examples of these platforms are Alt-Store, TutuApp, and AppValley, just to name a few. Unofficial sources for Android devices are Aptoide, APKPure, and F-Droid, among others. These third-party app stores allow you to download and install apps that are not available through the official channels. These apps can include modified versions of popular games, apps that have been removed from the Google Play Store, or Apple App Store, or even entirely new applications that haven't been approved by Google and Apple. While these sources can offer a wider variety of apps, they come with significant risks.

One of the primary concerns with using unofficial sources is the potential for security vulnerabilities. Apps from these sources do not undergo the same rigorous review process as those on the official Play Store and App Store. This means they could contain malware, spyware, or other malicious code that can compromise your device's security and steal your personal information. Additionally, these apps might not receive regular updates, leaving them susceptible to bugs and compatibility issues with newer versions of Android and iOS.

Despite the risks, some developers choose to distribute their apps through unofficial sources to reach a broader audience or to bypass the strict guidelines imposed by Google and Apple. For example, developers might use these platforms to test new features or gather user feedback before releasing an official version. However, you should exercise caution and thoroughly research any third-party app store and the apps they offer before downloading.

Sideloaded

Sideloaded apps involves installing software on your device from sources other than the official app stores, such as the Apple App Store or Google Play Store. This process bypasses the standard security and vetting procedures that official stores enforce, potentially exposing your device to unverified and potentially harmful applications.

The risks are identical to using a third-party store. While sideloading can offer access to a wider range of apps, including those not available in your region or apps that provide additional features, they also carry significant risks. These include the possibility of downloading malware, missing out on automatic updates, and compromising your privacy due to excessive permissions requested by unofficial apps. Therefore, while sideloading can be tempting for its flexibility, it requires careful consideration and caution to avoid compromising your device's security and performance.

An Android package (APK) is a developer file format for installation of Android applications. When developing an Android application, the developer will sideload the application, usually using an Android tool such as Android Debug Bridge (ADB). The ADB will allow the developer to install the APK directly onto the device they are testing with.

An iOS App Store Package (IPA) is a file format used for the installation of iOS applications. When developing an iOS application, developers typically use Xcode, Apple's integrated development environment, to build and test their apps. The IPA file can be sideloaded onto an iOS device using tools like Xcode, or third-party software such as Cydia Impactor. This process allows developers to install and test their applications directly on their iPhones or iPads before submitting them to the App Store.

Developer Mode

The *developer mode* on Android and Apple devices allows a developer to connect to the device via a USB connection. A developer can then create a bridge from a computer to side-load applications as well as debug. The developer mode on Android offers myriad settings that can be changed, such as viewing running services, staying awake, setting a mock GPS location, and USB debugging, just to name a few. The Android operating system allows you to change and tweak settings that can be security concerns. Because of all these tweaks and settings changes, the developer mode is a security concern. By using the mock GPS locations, you can make your phone think that it is within a geolocation security perimeter and possibly circumvent security.

Apple's latest iOS does not have a developer mode like Android. However, you can perform development functions through the Xcode application on a Mac to help you develop iOS applications.

The development mode on Android can be accessed by navigating to Settings > About Phone > Software Information, then tapping Build Information seven times. The Development menu will be on the parent menu under About Phone.

Root Access/Jailbreak

The terms *root access* and *jailbreak* are synonymous with each other. However, *root access* is normally associated with Android and *jailbreak* is normally associated with Apple iOS. When you attain root access for an Android device, you literally have access as the super user named root. This access allows you to change various aspects of the operating system, such as turning on premium features like hotspot and tethering, and changing the operating system on the device by flashing new firmware.

When you root an Android phone and flash a new firmware, you no longer have the patch management from the parent vendor of the phone. For example, if you root a Samsung phone and install Havoc OS, then you no longer receive the Samsung security policies and updates. The Google Play store will also be affected, since only older versions of the applications will be available for download. This means that patch management will lag behind.

When jailbreaking the Apple iOS, you attain a higher level of access to the iOS, just as you do when you root an Android device. The motivation to jailbreak an iOS device is to access premium features, such as adding photo modes, hiding apps, and installing newer features on old devices. Jailbreaking is a security concern because you are modifying the operating system of the phone. Malicious software can easily be installed, and with the new level of access, it can hide itself.

Many organizations that employ mobile device management (MDM) create policies to prevent rooted and jailbroken devices from attaining access to organization information. This policy might restrict getting email from the device, or it could restrict access to the network in your organization.

It is best to use the operating system that has shipped with your mobile device. In many cases, once a device has been rooted or jailbroken, it cannot be reverted back to stock.

Malicious Applications

A *malicious application* is any application with malicious intent for the user or the user's device. You can find malicious applications on both Android and Apple mobile platforms. You can identify a malicious app by reviewing its permissions and contrasting it with its function—for example, if you download a camera application and it asks to record calls. There should be some suspicion and you should revoke the permissions and potentially uninstall the software.

The obvious security concern for a malicious application is that it has excessive permissions to your data or your organization's data. Periodically you should review the permissions each application installation has on your device. Another rule of thumb is to look before you click and think of the security concerns before installing software.

Bootleg Applications

A *bootleg application* is a premium application that has been cracked or nullified to remove the digital rights management (DRM). Bootleg applications can be found for a number of premium mobile apps; they generally are in the form of an APK. Bootleg applications usually contain malicious software, because that is how the bootlegger makes their money. This obviously goes back to the discussion of verifying the source of the application and being cautious with APK installations.

Application Spoofing

Application spoofing is the act of a malicious application spoofing a legitimate application. Application spoofing is much more prevalent in the mobile application marketplace and can be observed on both Apple and Google mobile platforms.

The security concern for application spoofing is the possibility of installing malicious software on your device. Disguising itself as a legitimate application, it has the same security concern of access to personal data and organizational data. You can prevent application spoofing by verifying the name of the publisher, the icon for the application, and the number of installations. If you are downloading a social media platform and it has only 100 downloads, this should send up a red flag. Another method of validating the application is to read reviews for the application. Again, think before you click and download the software.

Common Symptoms

The following sections discuss common symptoms of problems with mobile operating systems and application-related security issues. As with so many issues involving troubleshooting, common sense is most important. Using logic and a systematic approach, you can often identify and correct small problems before they become large ones.

High Resource Utilization

High resource utilization can be a telltale sign that a device is running more than you think it should be—perhaps the drives are being searched or the camera is recording your every move. Monitor for high resource usage. If you discover it, find out what is causing it and

respond appropriately. In this section I will cover some basic components and what you should look out for.

High Network Traffic

A higher than normal amount of traffic can be a symptom of a security issue. Spikes in traffic for extended periods of time can mean that data is being stolen from your device or relayed through your device. You should have an idea of the volume of traffic you would normally expect on your device.

You can start closing applications as you watch the volume of traffic. When the volume of traffic subsides, you probably have your culprit. Then, as covered previously, check the application's permissions to see if the application is malicious or is compromised in some way. Clear cache and data, then uninstall and reinstall the application from a known good source. This procedure may identify the issue or verify that you had a malicious application installed. A telltale sign is if the application is no longer available.

Data-Usage Limit Notification

Exceeding the limits on data plans can also be symptomatic of a security issue. Data usage coincides with the volume of traffic previously discussed. A malicious application running on the device could be used to send *spam* or malware, or to conduct a multitude of other malicious activities from your device. A malicious application can also continually spy on you and your data. All of these activities can rob you of precious data in your data plan, pushing you over your contracted limits.

Excessive malicious use of data on a mobile device can be mitigated with two methods:

- The first method is watching the normal usage of data from month to month. Identifying a normal baseline of usage can alert you when data usage is abnormal.
- The second method is to use a mobile firewall, which limits the traffic leaving the mobile device.

Sluggish Performance

While applications, normal usage, and so on can contribute to *degraded performance*, another offender could be malware or a virus. When you observe degraded response time on your device, you need to investigate the symptom, as it could indicate a security issue. Check RAM and CPU usage, and if an application is out of control, it could be infected with malware. It is best to run an antivirus/antimalware scan on the device to check it thoroughly.

Not every problem is related to a possible security threat. The normal search for a cellular signal can be just as taxing on the device. However, if you are in the normal locations in your day-to-day travels, such as work and home, and still experience degraded performance, you may have an application problem (out of memory) or a security threat. In either case, the issue needs to be checked out quickly.

Limited Internet Connectivity

When you have limited Internet connectivity on your mobile device, you should not immediately think that the limited connectivity is a security symptom. Limited Internet connectivity can be a result of many different problems. Mobile devices are very susceptible to limited

Internet connectivity because they contain small transceivers for wireless and cellular communications. The radio firmware also plays a big role in choosing the right radio frequency and is often the problem with connectivity.

Taking everything into consideration, if you know nothing has changed in the wireless environment, your firmware is the same, you've rebooted, and no one else is having a problem, you can suspect this is a security symptom. Malicious applications will often monopolize your connection or proxy the connection in an effort to sniff usernames and passwords. Both monopolization of the connection and proxying of the connection can create intermittent Internet connections. If your connection is being monopolized, then you will see high network bandwidth as previously discussed. There are a number of ways that a network connection can be proxied for malicious purposes, such as DNS proxy, network transmission, and wireless, just to name a few. The ideal way to combat this issue is with a mobile device firewall and antimalware software.

No Internet Connectivity

If all the usual causes have been reviewed, then you should suspect no Internet connectivity to be a security-related symptom. In some rare instances, malware will cause the mobile device to have no Internet connectivity. This generally happens because a DNS server that the mobile device is pointed to for malicious reasons has ceased to function or the relay server that is proxying the connection has ceased to function. There are a multitude of reasons why no Internet connectivity would be experienced. As previously recommended, a good mobile device firewall and antimalware should be employed.

Complete failure of an Internet connection is easier to diagnose than an intermittent problem. Therefore, a factory reset of the device should aid in figuring out whether the hardware is bad or the problem is software, such as malware, triggering the problem.

High Number of Ads

When a mobile device is experiencing a high number of ads, this is a security-related symptom of adware. Adware is a type of malware that pops up ads for malicious purposes, usually to entice the user to buy something. Adware is usually the result of installing a malicious application on the mobile device. There are two ways to diagnose the problem of adware. The first depends on the number of ads and the frequency of the ads. You can start by uninstalling applications until the ads stop popping up. This method is preferable because it is probably the quickest.

The second method, a factory reset, is much more effective, but it will not identify the malicious application. A factory reset will remove any malware along with all the applications. Of course, you can start installing applications until the ads start. However, resetting a phone is a pretty anxiety-producing process for avid mobile device users frantically trying to sign back into applications. Therefore, the first method may be preferred. Once the application is uninstalled and the device is factory reset, be sure to install antimalware software on the device.

Fake Security Warnings

A *fake security warning* on any system is a big red flag and a security symptom. Mobile devices are not exempt from fake security warnings, although these warnings on full operating systems are more common. Regardless, when a fake security warning is

discovered, it should be treated as if malware is installed on the device. The device should be factory reset and antimalware should be installed prior to reinstalling the applications.

Unexpected Application Behavior

Unexpected application behavior is not always an indication that you have been infected with malware or have a security symptom. Applications have unexpected behavior all the time. However, when an untrusted and newly installed application behaves in an unexpected way, this could be a symptom of a security problem with the application.

When you experience unexpected application behavior, you should immediately question the trust of the application. This can be done by reading reviews for the application to determine if others have run into similar problems. Also judge the application by its installer base, which is proportional to the reviews. For example, if an application has 100 installs and only has 5 people commenting that it's a great application, then this app should fall under suspicion.

The first step to be taken is to scan the device for malware. If the application is flagged as malware, then a factory reset should be performed. Then install only the trusted applications that you use daily.

Leaked Personal Files/Data

When authorized users access devices through unintended connections or unauthorized users access stolen devices, they can access the data on the device. Outside of these risks, there is always the risk of loss or theft of the device itself.

Therefore, security for mobile devices should be applied in a layered approach. Antivirus and antimalware software should be installed on the device to protect it from malicious applications. In addition, a mobile firewall should be installed along with the antivirus and antimalware software. Fortunately, there are third-party security suites that can protect you from all these threats.

Mobile device management (MDM) software should also be employed. This software is like the Swiss army knife of security for mobile devices. It can require passcodes, the installation of antivirus, antimalware software, mobile firewalls, current updates, and so much more. One of the most notable features is the ability to remotely wipe the device in the event it is stolen or lost.

In addition, there should be a firm policy that details the encryption of data in use, at rest, and in transit. A written policy should be drafted along with procedures on how to deal with leaks when they occur. These policies are usually drafted with an insurance company in order to protect an organization in case of a data leak of personal information.

Summary

This chapter addressed systematic approaches to working with computer problems as well as troubleshooting operating systems and resolving security-related issues. In our discussion of troubleshooting theory, you learned that you need to take a systematic approach to problem solving. Both art and science are involved, and experience in troubleshooting is helpful but not a prerequisite to being a good troubleshooter. You learned that in troubleshooting, the first objective is to identify the problem. Many times, this can be the most time-consuming task.

Once you've identified the problem, you need to establish a theory of why the problem is happening, test your theory, establish a plan of action, verify full functionality, and then document your work. Documentation is frequently the most overlooked aspect of working with computers, but it's an absolutely critical step.

Next, we discussed operating system–related troubleshooting issues. First, we looked at common symptoms, and then we discussed some tools that can be helpful in solving problems.

Finally, we looked at security-related troubleshooting issues as well as best practices for removing malware. Again, we started by looking at common issues and then how to solve them.

Exam Essentials

Know the common symptoms and solutions for operating system problems. Be able to identify the five critical systems of CPU, RAM, network, disk, and graphics that interact with an operating system. Know the various symptoms that will lead to operating system problems. Also, be able to identify solutions to operating system problems.

Understand the various security issues related to operating systems. Be able to identify common security issues that you will see affecting operating systems. You should also understand the impact of these security issues on both the operating system and the organization, as well as common ways to mitigate and solve these security issues.

Know the process for removing malware. Know the seven steps to remove malware. Be able to identify how to perform these seven steps during the malware-removal process. You should be familiar with the various tools used to identify and remove malware.

Know how to troubleshoot mobile OS and application issues. Be able to identify the common symptoms of application-based issues for mobile devices. Know how to solve mobile issues for applications on both Android and Apple devices.

Understand mobile device security issues. You should have a good understanding of the various mobile device security symptoms, problems, and solutions. Be able to articulate how these issues pose a risk to the organization, as well as how these issues can be mitigated and solved.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. In Windows, which utility is responsible for finding, downloading, and installing Windows patches?
 - A. Device Manager
 - B. Microsoft Management Console
 - C. Download Manager
 - D. Windows Update
2. Which Startup Setting option allows you to boot with basic drivers?
 - A. Enable Debugging
 - B. Enable Safe Boot
 - C. Disable Driver Signature Enforcement
 - D. Enable Low-Resolution Video
3. Which bootrec option can be used in Windows to rebuild the boot configuration file?
 - A. /fixboot
 - B. /rebuildbcd
 - C. /scanos
 - D. /fixmbr
4. What is the first step in malware removal?
 - A. Quarantine the infected system.
 - B. Identify and verify the malware symptoms.
 - C. Remediate the infected system.
 - D. Educate the end user.
5. Which tool will allow you to troubleshoot a slow-loading profile?
 - A. Profile tab of the Advanced System Properties
 - B. Regedit
 - C. Windows Recovery Environment
 - D. Windows Preinstallation Environment
6. Which of the following components are only used to restore Windows from a suspended state?
 - A. BCD
 - B. ntoskrnl.exe
 - C. winload.efi
 - D. winresume.exe

7. One of the users you support has a Windows 11 laptop that will not boot up. The user just installed brand-new drivers for a graphics card. They need to access a tax application and their data files. What should you try first?
 - A. Use System Restore.
 - B. Use Reset This PC.
 - C. Reimage the laptop.
 - D. Manually reinstall Windows 10.
8. Which partitioning type is required when you have UEFI firmware?
 - A. GPT
 - B. MBR
 - C. POST
 - D. Boot Sector
9. Which of the following are used to prevent pop-up windows from appearing?
 - A. Antimalware utilities
 - B. Pop-up blockers
 - C. Phishing sites
 - D. Antivirus software
10. In general, how often should you update your antivirus definitions?
 - A. Weekly.
 - B. Monthly.
 - C. Daily.
 - D. Antivirus definitions do not need to be updated.
11. Which tool can be used to diagnose why Windows 11 is slow and sluggish?
 - A. Resource Monitor
 - B. msconfig.exe
 - C. Device Manager
 - D. Reliability Monitor
12. Which tool will allow you to diagnose why Windows Update keeps failing?
 - A. ntbtlog.txt
 - B. Windows Update Troubleshooter
 - C. Windows Recovery Environment
 - D. Safe mode
13. Which of the following programs could be considered antimalware?
 - A. Windows Defender Security
 - B. MDM
 - C. Windows Action Center
 - D. VirusTotal

14. Which of the following tools allows you to manually fix maliciously modified system files?
 - A. regedit
 - B. SFC
 - C. bootrec
 - D. UAC
15. Which of the following can you do to help eliminate security problems? (Choose the best answer.)
 - A. Establish security policies and procedures.
 - B. Optimize drives.
 - C. Prevent booting into Safe Mode.
 - D. Prevent booting into Windows Recovery Environment.
16. A mobile device is running out of RAM. What could be the most likely problem?
 - A. The device is not charged to capacity.
 - B. The digitizer is not functioning properly.
 - C. The device is in DND mode.
 - D. The device has background applications open.
17. What is a risk of using the auto-reconnect feature on a mobile device?
 - A. The device will reconnect to any SSID.
 - B. The device could be exploited by an evil twin attack.
 - C. The device's battery life could be shortened.
 - D. You may exceed your cellular data plan's limits.
18. You notice that the reliability of the operating system has diminished in Reliability Monitor. Where can you find more details on why applications are failing?
 - A. Device Manager
 - B. Event Viewer
 - C. Windows Recovery Environment
 - D. msconfig.exe
19. Why would the operating system write out large amounts of RAM to the page file?
 - A. The CPU is running high on utilization.
 - B. This is a normal process of the operating system.
 - C. The amount of physical RAM is low.
 - D. The page file is faster than conventional RAM.
20. What is one consequence of an overheating mobile device?
 - A. Higher RAM usage
 - B. Degraded battery life
 - C. Inaccurate touchscreen response
 - D. Inability to decrypt emails

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

List, in order, the seven best practice steps associated with malware removal.



Chapter 8

Scripting, Remote Access, and Artificial Intelligence

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **4.0 Operational Procedures**
- ✓ **4.8 Explain the basics of scripting.**
 - Script file types
 - .bat
 - .ps1
 - .vbs
 - .sh
 - .js
 - .py
 - Use cases for scripting
 - Basic automation
 - Restarting machines
 - Remapping network drives
 - Installation of applications
 - Automated backups
 - Gathering of information/data
 - Initiating updates
 - Other considerations when using scripts
 - Unintentionally introducing malware
 - Inadvertently changing system settings
 - Browser or system crashes due to mishandling of resources



✓ **4.9 Given a scenario, use remote access technologies.**

- Methods/tools
 - RDP
 - VPN
 - Virtual network computer (VNC)
 - Secure Shell (SSH)
 - Remote monitoring and management (RMM)
 - Simple Protocol for Independent Computing Environments (SPICE)
 - Windows Remote Management (WinRM)
 - Third-party tools
 - Screen-sharing software
 - Videoconferencing software
 - File transfer software
 - Desktop management software
- Security considerations of each access method

✓ **4.10 Explain basic concepts related to artificial intelligence (AI).**

- Application integration
- Policy
 - Appropriate use
 - Plagiarism
- Limitations
 - Bias
 - Hallucinations
 - Accuracy
- Private vs. public
 - Data security
 - Data source
 - Data privacy



CompTIA has identified that with the rapid adoption of cloud-based services, system administrators have a need for scripting and remote access now more than ever. Scripting allows you to manage a system as if you have logged in and performed the task yourself, such as purging old log files from a web server to make space. Remote access technologies allow you to connect to these remote systems so that you can manually administer them.

This chapter looks at a variety of scripting languages. After completing this chapter, you may not be an expert at writing scripts—the goal is to familiarize you with the various scripting languages so that you can learn the basics of scripting and their purposes. We'll also look at various remote access methods and their security considerations.

A new addition to the CompTIA objectives is artificial intelligence (AI). The subject is an inch deep and 5 miles wide. So, we will cover some a myriad of topics that will make you more knowledgeable with AI, while covering the objectives for the exam.




Exam Objective 4.8, “Explain the basics of scripting,” requires you to be able to identify and read basic script types of Windows batch scripts, PowerShell, Visual Basic Script, Bash, JavaScript, and Python. You should also be able to apply scripting concepts to use cases. In addition, you should know the considerations when scripting.


Scripting

Before diving into scripting, we'll begin by discussing the differences between a programming language and a scripting language. When you write a script, you are basically using a high-level programming language. An example of a low-level programming language is assembler, also known as *assembly language*. Figure 8.1 shows what we call the “programming pyramid.”


The lowest layer is the actual hardware, which is the central processing unit (CPU). Directly above the hardware layer is the *executable machine code* that interacts with the hardware to make it perform some useful function or process. The operating system is usually programmed in a low- to mid-level programming language, such as C/C++ or even assembly language. It is then *compiled* into executable machine code. Applications are often programmed in high-level languages, such as Java, C#, or VB.NET, and are compiled to executable machine code or an intermediate code. Scripts, however, are not compiled; they are interpreted, as we discuss in this section.

FIGURE 8.1 Programming pyramid

Depending on which layer you program on, you gain some advantages, but at the same time you are also capped by some of the limitations of each layer. For example, a program created in assembly language will be quite complex because you will need to perform low-level functions just to add two numbers together. However, because you are at such a low level, you have direct access to the hardware, so you are limitless in your control of the hardware. If you were to use a higher-level language, such as C/C++, the application would be relatively easy to write; adding two numbers is as simple as `c = a + b;`. Because it is a higher-level language, however, you do not have the same lower-level control over the hardware as with assembly language. You have access only to what the compiler will understand and compile to executable machine code, as shown in Figure 8.2.

FIGURE 8.2 Compiling a programming language

Scripting languages do not need to be compiled. They are interpreted by the shell, command line, or external interpreter, as shown in Figure 8.3. The interpreter reads the script and executes the instructions in the operating system. The big difference is that you do not need to compile scripts to executable machine code, as in the previous example of a C/C++ program. Unfortunately, the higher the level, the less control you have over the process. The benefit is you can create a script rather quickly, and you don't need to compile it.

FIGURE 8.3 Interpreting a scripting language

Another big difference between applications and scripts is that scripts require applications to complete their purpose. If an application doesn't exist for a function in your script, then you should evaluate whether a script is the right course of action. For example, if you need to resize pictures, an application must exist that can be called via the command line that will take the appropriate input and resize the pictures. If an application doesn't exist that can resize the picture, then you might need to write an application that resizes a picture in lieu of a script. A script cannot normally create functions of this nature; it can only call on them.

**NOTE**

Java is a compiled language that requires the Java Development Kit (JDK) in order to be compiled to an intermediate code called bytecode. The bytecode is then run on top of the Java runtime, known as the Java Runtime Environment (JRE). This is why you need to install Java to run Java programs on your operating system. JavaScript, on the other hand, is a scripting language that many different web browsers support. JavaScript does not need a JRE since it is interpreted inside the web browser.

Scripting Basics

There are several different types of scripting languages. Each scripting language has its own nuances and syntax. This section discusses some scripting basics that you can apply to any scripting language. You'll only have to adjust the syntax to the scripting language you've selected to get the desired outcome.

Variables

A *variable* is a symbolic word or combination of letters that can be used to hold a value. The value can be defined as a number, text, Boolean, or array. Numbers can be either *integers* or *floating-point*, but the scripting languages you choose must support floating-point math. Floating-point math is a value that contains a decimal number, whereas integers are whole numbers. Text is also called a *string* and has no numeric equivalent. A string containing 30

will not be a value of 30; it will be only be a string of characters containing a 3 and a 0. Mathematical computations cannot be performed on a string. Boolean values are true or false values, and true values and arrays are collections of strings, numbers, and Boolean values.

You might see a programming language described as a strongly typed language. That means that the variable must be defined as to what type of value it will hold and how long the value will be. This is yet another difference between scripting languages and programming languages: scripting languages do not need variables to be defined and are normally not typed. Most of the time, you can just load a variable with a value and it is dynamically typed. This is not the most economical use of memory, but that shouldn't be an issue, because scripts are often simple ways to complete a simple task; therefore, it doesn't matter that loading a variable with a value and having it dynamically typed isn't the most economical use of memory.

The name of a variable should have some meaning and cannot be a reserved word. So if you had a variable that needed to hold a counter, you could simply create a variable called `count`. However, you would need to ensure that `count` was not used elsewhere in the program for a function; otherwise, it's a *reserved word* in the scripting language and cannot be used as a variable.

The following is an example of a variable named `count` being created in the *PowerShell* scripting language and loaded with the value 1. Generally, you load a value by specifying the variable, followed by an equal sign, and then the value. In PowerShell, the syntax dictates that you place a `$` in front of the variable in order to notify the interpreter that you are addressing a variable. Each scripting language is slightly different, but the concept remains the same. You can also view the contents of the variable either by typing the variable name and pressing Enter or by using the `echo` command before the variable.

```
PS C:\Users> $count = 1
PS C:\Users> $count
1
PS C:\Users> echo $count
1
PS C:\Users>
```

Environment Variables

As previously mentioned, variables hold values in scripts so that we can do things such as counting. Environment variables also hold values, but they are used for the environment of either the system or the current user. Environment variables often hold values like the path to an executable or the location of the temporary folder.

Environment variables are inherited in a structured fashion: system, user, and then program. They can be overwritten at any underlying level, but only for that entity. For example, if a script changes the user environment variable for the temporary folder, the change is applicable only for that script unless the user environment is changed—then it will affect all scripts run by a user. Here are the types of variables you will encounter in the operating system:

System Variable Defined for use by the system, or any user or program

User Variable Defined for use by a specific user, or any program executing under that specific user

Program Variable Defined for use inside a script or application and can only be accessed from inside the script or application

Comment Syntax

Your scripts should have a level of end-user readability. The code should describe what is going to happen without external documentation. Although we strive to create eloquent scripts that read like computer poetry, sometimes we need to make comments inside the code.

Each language has a different way to make comments. The following example uses the PowerShell syntax. As you progress through the next section, “Scripting Languages,” you will see the way to comment in the various languages. You’ll notice in the following code that we’ve added some commenting that is generally used for script creation. The three basic components of the comments are author name, authored date, and the purpose of the script. Although the following comment block is best practice, it is not required; one-line comments inside the code are allowed also. In this scripting language, PowerShell, the comment is preceded by the # symbol. Other scripting languages have their own syntax, and you’ll see the various examples throughout this chapter.

```
# Sybex
# 4/20/25
# This script will output a directory listing of the C drive
Get-ChildItem C:\
```

Loops

When writing a script, you may need to create a controlled loop inside the script. An example of a controlled loop is to read a file and then do some things for each line contained in the file. You can use two basic loops in your scripts: `for` loops and `while` loops.

A `for` loop is a stepped loop with a defined beginning and a defined end; each step is defined as well. The following example shows a `for` loop for PowerShell that starts at 1 and counts to 10. The output is the word *Number* and an incrementing number. You can see that the variable is initially set to 1 with (`$count=1`), and a test is done to check if the variable is less than or equal to 10 with (`$count -le 10`). The `for` loop is stepped by adding 1 to the existing number with (`$count++`). This is just one example of a `for` loop, and every scripting language is a bit different in syntax, but the concept is the same.

```
For($count=1; $count -le 10; $count++) {
    Write-Host "Number $count"
}
```

[Output of Script]

Number 1

```
Number 2  
Number 3  
Number 4  
Number 5  
Number 6  
Number 7  
Number 8  
Number 9  
Number 10
```

A `while` loop continues to loop until either it is exited or a condition is met. The `while` loop has no defined beginning, only a defined end, and it can be exited without consequence to the function. In the following code, the variable `count` is set to 0, and then the `while` loop begins. Inside the `while` loop, the `count` variable is incremented by 1 each time it loops with (`$count++`). The loop will continue as long as `count` is less than 11 with (`$count -lt 11`). The output will be identical to the prior example.

```
$count = 1  
While ($count -lt 11) {  
    Write-Host "Number $count"  
    $count++  
}
```

[Output of Script]

```
Number 1  
Number 2  
Number 3  
Number 4  
Number 5  
Number 6  
Number 7  
Number 8  
Number 9  
Number 10
```



A `while` loop and a `do while` loop are identical in functionality; they both loop until a condition is no longer true. The difference between the two is that a `while` loop puts the condition before the code in the loop, whereas a `do while` has the condition after the code in the loop. Each language supports either one loop or the other, and sometimes both types of loops. However, the concept remains the same: The loop will continue until the condition is no longer true.

Branch Logic

Branch logic enables code to deviate, or branch, depending on a condition. The `if` statement is the most common conditional branch logic found in scripts. The `if` statement is usually followed by a condition and a then clause. It can even use an `else` clause. The then clause is executed only if the condition is true; if the condition is not true, the `else` clause is executed.

In the following example, the string `str` is set to 1. The `if` statement then checks the condition of `str` being equal to 1 with `($str -eq 1)`. It is true, so the next statement of `Write-Host "Yes"` is executed. It's important to note that PowerShell implies the then clause; only the `else` clause needs to be spelled out. If `str` were anything other than 1, the condition would be false, executing the `else` clause of `Write-Host "No"`. This is an example for PowerShell, but in the next section, "Scripting Languages," you will find many different examples for each scripting language.

```
$str=1
If ($str -eq 1) { Write-Host "Yes" } Else { Write-Host "No" }
[ Output of Script ]
Yes
```

Scripting Languages

Now that you understand the basic differences between a programming language and a scripting language and some scripting basics, let's look at several different scripting languages you are likely to encounter as a technician. As we cover the various languages, we'll highlight which operating systems they are common to or natively supported on—which is often the biggest deciding factor when you choose a scripting language.

Windows Batch Scripts

Windows batch scripts have been around since the release of Microsoft's Disk Operating System (DOS) back in 1981. The original script interpreter was `command.com`, and since then, Windows NT was released, which included an updated command-line interpreter called *Command Prompt*, or `cmd.exe`. The original file extension used with batch scripts was `.bat`, but today both `.bat` and `.cmd` can be used to initiate a batch script because they are both associated with the command-line interpreter `cmd.exe`.

A Windows batch script is probably the fastest way to get something done when all you need is a list of commands run one after the other. For example, say you need to create new user accounts for a school. You can get the usernames in an Excel sheet, and you can create a script from the entries by using an Excel formula, adding in the column of the username. As an example, the formula `= "NET USER" & A1 & "PassW0rd /ADD /DOMAIN"` copied into cell B1 will produce the line you need to execute. Then you just need to drag the formula down, and the script will be built. A quick copy and paste, and the script will look similar to the following output. It's a quick and dirty way to create user accounts. Using the

combination of an Excel sheet and a copy-and-paste into a batch script is the fastest way to build and execute a laundry list of commands.

```
NET USER UserOne PassW0rd /ADD /DOMAIN
NET USER UserOne PassW0rd /ADD /DOMAIN
NET USER UserOne PassW0rd /ADD /DOMAIN
NET USER UserOne PassW0rd /ADD /DOMAIN
[ Output Cut ]
```

Batch scripts can also contain logic. The following is a simple batch script that tests whether a variable of FLIPFLOP is equal to 0 using an `if` statement. If FLIPFLOP is equal to 0, the script will proceed to write to the screen the word *Zero*, set FLIPFLOP to a value of 1, and then jump to :LOOP. Because FLIPFLOP is set to 1, it is not equal to 0; so, the `else` clause in the `if` statement will be processed, and the word *One* will be printed to the screen, and FLIPFLOP will then be set to 0, and the script will jump to :LOOP again. This will proceed until Ctrl+C is pressed to stop the processing of the script.

```
@ECHO OFF
REM FlipFlop Script
SET /A FLIPFLOP=0
:LOOP
IF %FLIPFLOP% EQU 0 (ECHO Zero && SET /A FLIPFLOP=1) ELSE (ECHO One && SET /A
FLIPFLOP=0)
GOTO :LOOP
```

In Exercise 8.1 you will create a simple Windows batch script from the preceding script. The end result will print out *Zero* and *One* in an alternating pattern.

EXERCISE 8.1

Creating and Running a Windows Batch Script

1. Click the Start menu and type **Notepad**. Select Notepad from the search results.
2. Type the preceding batch script example into Notepad.
3. Choose File ➤ Save As, and then save the batch script to your desktop.
4. Name the file **flipflop.cmd**, and then click Save.
5. Press the Windows key and the R key at the same time, type **cmd.exe**, and then press Enter.
6. Type **cd %userprofile%\desktop** and press Enter.
7. Type **flipflop.cmd**, and then press Enter.

Notice that your screen display will alternate between *One* and *Zero*.

8. Press Ctrl+C.

This will display the question `Terminate batch job (Y/N)?`.

9. Answer Y, and then press Enter.
10. Switch to Notepad and add **REM** in front of the `@echo off` command.
11. Save the file by choosing File > Save.
12. Switch to the command prompt and rerun `flipflop.cmd`.

You will now see the commands as well as the variable contents while the script is running.

13. Press Ctrl+C.

This will display the question `Terminate batch job (Y/N)?`.

14. Answer Y, and then press Enter.
-



Another way to open a command prompt to a specific folder is to first browse to the folder, and then hold the Shift key and right-click an empty space in the folder. Select `Open PowerShell Window Here` on Windows 11. After opening the PowerShell window, type **cmd** and then press Enter. This will change your interpreter to `cmd.exe`.

PowerShell Scripts

PowerShell allows for the automation and management of the Windows operating systems, as well as cloud-related services such as Microsoft Azure and Microsoft 365. One of the limiting features of any scripting language is its ability to perform a needed task. PowerShell was created to be totally extensible. It was built on the .NET Framework *Common Language Runtime (CLR)*. Any programmable library a .NET application has access to, PowerShell can use, which is what makes it so extensible. It has been used since Windows Server 2008 as a configuration tool for the operating system. In fact, most of the time when you configure a service in the Server Manager tool, you actually run a PowerShell command in the background. Many of the GUI wizards allow you to see the PowerShell script that will be executed so that you can reuse the line in a script of your own.

PowerShell introduced the concept of *cmdlets*. PowerShell has over 100 cmdlets installed, called the core cmdlets. You can always add your own cmdlet by creating a PS1 script and installing it into the PowerShell cmdlet store in the operating system, as you will do in Activity 8.2. A cmdlet is simply a verb and a noun separated by a dash. Here are a few examples:

Get-Item Gets an item such as a directory listing, environment variable, or Registry key

Set-Item Changes the value of an item, such as creating an alias or setting an environment variable

Copy-Item Copies an item, such as a file or folder

Remove-Item Deletes an item, such as a file, folder, or Registry key

Move-Item Moves an item, such as a file, folder, or Registry key

These are just a few of the built-in core cmdlets for PowerShell. Others exist for `-Item`, such as `Rename-Item`, `New-Item`, `Invoke-Item`, and `Clear-Item`. Each one performs a corresponding action on the noun following the dash. You can even extend the functionality of a command with your own PS1 cmdlet. There is a `Get-Verb` command so that you can see all the appropriate verbs that you can use for your own command.



To learn more about creating your own cmdlets, visit:

<https://docs.microsoft.com/en-us/powershell/scripting/developer/cmdlet/how-to-write-a-simple-cmdlet?view=powershell-7.2>

If you use the `Get-Item` cmdlet and specify a folder, information about that folder will be returned. If you want to see all the other folders contained within that folder, you can use a `*` wildcard. Or you can use the `Get-ChildItem` cmdlet and specify the directory, as follows:

```
PS C:\Users\UserOne> Get-item c:\*
```

```
Directory: C:\
```

Mode	LastWriteTime	Length	Name
----	-----	-----	----
d-----	11/28/2024 9:52 PM		Dell
d-----	5/16/2024 9:32 PM		NVIDIA
d-----	4/11/2024 7:38 PM		PerfLogs
d-r---	5/28/2024 10:04 PM		Program Files
d-r---	8/12/2024 5:31 PM		Program Files (x86)
d-r---	5/28/2024 6:11 PM		Users
d-----	10/18/2024 10:17 PM		Windows

```
PS C:\Users\UserOne> Get-ChildItem c:\
```

```
Directory: C:\
```

Mode	LastWriteTime	Length	Name
----	-----	-----	----
d-----	11/28/2024 9:52 PM		dell
d-----	5/16/2024 9:32 PM		NVIDIA

d-----	4/11/2024	7:38 PM	PerfLogs
d-r---	5/28/2024	10:04 PM	Program Files
d-r---	8/12/2024	5:31 PM	Program Files (x86)
d-r---	5/28/2024	6:11 PM	Users
d-----	10/18/2024	10:17 PM	Windows

When you use the `dir` command in PowerShell to view a directory listing of files, you are actually using something called an *alias*. The alias then calls the `Get-ChildItem` cmdlet. To see all the aliases on the operating system, you can use the `Get-Alias` cmdlet. You can see all the commands mapped over to PowerShell cmdlets, as follows:

```
PS C:\Users\UserOne> get-alias
CommandType      Name
-----
Alias             % -> ForEach-Object
Alias             ? -> Where-Object
Alias             ac -> Add-Content
Alias             asnp -> Add-PSSnapin
Alias             cat -> Get-Content
Alias             cd -> Set-Location
Alias             CFS -> ConvertFrom-String
Alias             chdir -> Set-Location
Alias             clc -> Clear-Content
Alias             clear -> Clear-Host
Alias             clhy -> Clear-History
Alias             cli -> Clear-Item
Alias             clp -> Clear-ItemProperty
Alias             cls -> Clear-Host
[ Output Cut ]
```

PowerShell also has a great way to develop scripts in what is called an *Integrated Scripting Environment (ISE)*. The ISE allows you to write a script and test it without having to switch back and forth between a text editor and the execution environment. In addition to the writing and execution environment in the same window, there is a type-ahead feature that allows you to pick a command if you remember only the first few letters, as shown in Figure 8.4. You can also use the Tab key to complete a command, which makes writing scripts easy when you know the first couple of letters. Formatting is also automated and makes for easy-to-read scripts. The formatting highlights variables and commands in different colors so that you can differentiate between the two.

Before any script can be executed on the Windows operating system, you must first allow scripts to run. By default, any PowerShell scripts will be blocked. Exercise 8.2 shows you how to “unrestrict” PowerShell scripts using the `Set-ExecutionPolicy` cmdlet.


```
Write-Host "One"  
$ FLIPFLOP=0  
}  
}
```

7. Click File ➤ Save.

The save dialog box will default to your Documents folder.

8. Rename the file **flipflop.ps1**, and then click Save.
9. In the lower portion of the ISE, type **cd ~\Documents**, and then press Enter.

Note that you used the ~ character, which represents your home directory, as previously covered in Volume 2 Chapter 4, “Working with macOS and Linux.”

10. Type **.\flipflop.ps1**, and then press Enter.

This will run the script and display the alternating words *Zero* and *One*. To stop it, press Ctrl+C. Alternatively, you can click the red stop button on the ribbon at the top of the window. You can execute the script by clicking the green play button.

11. Type **set-executionpolicy default**, and then press Enter to change the execution policy back to the default setting.



Microsoft has created a cross-platform version of PowerShell, originally called *PowerShell Core*, that runs on Linux and macOS. This version will run nearly all the commands that the Windows version of PowerShell offers, with a few platform-specific nuances. You can find more information on this topic at <https://learn.microsoft.com/en-us/powershell>.

Visual Basic Scripts

Visual Basic scripts—also known as *VBScripts*—are scripts based on Microsoft’s Visual Basic language. VBScript technology has been around since Windows 98 and Windows NT, so it is very mature, but it is slowly being replaced by PowerShell, as we’ll explain. VBScripts also run only on Windows operating systems, unlike PowerShell scripts.

VBScript technology is based on the *Component Object Model (COM)* to allow interaction with the operating system. Any COM object that can be instantiated can be accessed through a VBScript. As you learned in the preceding section, PowerShell is based on the .NET Framework. The Component Object Model predates the .NET Framework as a way to register programming libraries with the operating system. By default, VBScripts cannot access programming libraries that have been written for the .NET Framework, which is why VBScripts are slowly losing popularity and support by Microsoft.

VBScripts are still extremely useful when it comes to a structured language for creating login scripts. As previously mentioned, scripts are an interpreted language and are not compiled. VBScripts are no different; they require an interpreter to process. There are three main interpreters that can process VBScripts: *Windows Scripting Host (WSH)*, Internet Information Services (IIS) Active Server Pages (ASP), and Internet Explorer. ASP and Internet Explorer are deprecated, so we will focus on WSH.

The Windows Scripting Host is an environment that allows you to run VBScripts from the command line. By default, when a VBS script is run, a program called `wscript.exe` processes the script. Any output will be sent to a Windows message box that you must close by clicking OK. This can be quite annoying if you have multiple lines of output, as each line will pop up a message box you have to close. A VBScript can also be executed with the `cscript.exe` program. This version of the VBScript processor outputs to a console window—the name stands for console script.

VBScripts are normally edited in Notepad and saved with a `.vbs` extension. A number of third-party editors allow type-ahead. As previously mentioned, type-ahead is a feature that allows for the completion of code after a few letters have been entered. Microsoft supports an editor called Visual Studio Code, which you can download from <https://code.visualstudio.com>.

The following VBScript will perform the same output as the two preceding examples. As you can tell, the syntax is a little different.

```
'FlipFlop Script
Do Until FLIPFLOP > 2
  If FLIPFLOP = 0 Then
    Wscript.Echo "Zero"
    FLIPFLOP = 1
  Else
    Wscript.Echo "One"
    FLIPFLOP = 0
  End If
Loop
```



Because of the popularity of Windows PowerShell and its access to the .NET Framework, there will be no future releases of the VBScript processor. Microsoft will maintain the VBScript processor until it is retired. However, Windows batch scripts would probably be deprecated before the VBScript processor.

Shell Scripts

Bash stands for the Bourne Again Shell. Bash is backward compatible with its predecessor, called the *Bourne shell* (or *sh*). Whether you're using the Bash shell or the *sh* shell, either will perform similar functions for scripting purposes. The shell itself interprets commands

for scripting, similar to the Windows DOS shell. However, the Bash and sh shells are much more advanced in their structure and functionality than DOS. Linux/UNIX shell scripts are basically how the operating system boots itself up and starts all the necessary services.

Linux and UNIX scripts often end with the `.sh` extension to signify to the end user that the file contains a script (text) and not an executable code. Shell scripts do not require the `.sh` extension in order to be executed. They do, however, require execute permissions to be applied. This can be done with the `chmod` command, as you learned in Volume 2 Chapter 4. Another nuance is that you must specify the relative path using `./` and then the script—for example, `./script.sh`.

Shell scripts can be edited in any text-based editor. The vi editor was probably the first editor—and on older systems, your only choice. Today, there is a multitude of editors from which to choose, such as pico/nano, jed, Gedit, and Kate/Kwrite, just to name a few. Any of these editors will do a fine job for the basic editing of scripts. Some people find that opening two consoles is the fastest way to write a script on Linux/UNIX. One console serves as the editing console, with the text editor loaded and saving changes that are made, while the other console serves as the execution environment.

The following is a simple shell script that performs identical functionality to the scripts you have seen throughout this chapter. There are two major things to note with this example. The first is the syntax of the script, which differs slightly from the other examples but basically looks similar. The second thing to note is the directive on the first line. The directive tells the operating system which shell to process the script in. The `#!` is called a *hashbang*. The hashbang lets the operating system know the path of the script interpreter that will be used to process the script—in this case, it is `/bin/bash`.

```
#!/bin/bash
#FlipFlop Script
while [ ! $FLIPFLOP ]; do
    if [ $FLIPFLOP -eq 0 ]
    then
        echo Zero
        FLIPFLOP =1
    else
        echo One
        FLIPFLOP =0
    fi
done
```



The best way to start scripting with Linux/UNIX shell scripts is to load up Linux and try it out. You can find a great guide to Linux scripting here:

<https://help.ubuntu.com/community/Beginners/BashScripting>

Python Scripts


The *Python* scripting language was first released in 1991, and it has started to gain popularity over the past 10 years or so. It is not normally installed by default in any operating system. If you want to use Python on a Windows or macOS operating system, you need to visit www.python.org to download the latest version, and then install it. If you are running a Linux operating system, you will install Python through the package management system of the operating system.



If you install Python from www.python.org, you will only be prompted to download Python 3.x. However, when installing on Linux, if you do not specify Python 3, you will install the older version, Python 2.x. Python 2 is no longer being supported and has fewer features than the current version, Python 3.x.

The Python installation does not include an *integrated development environment (IDE)*, also known as an *Integrated Scripting Environment (ISE)*. The installation contains only the interpreter and some documentation to get you started. You can use any text editor to create and edit scripts. An extremely popular Python IDE is the Free Community version of PyCharm, by JetBrains. The IDE allows for script development similar to the PowerShell ISE but with many more features, as shown in Figure 8.5.

FIGURE 8.5 The PyCharm IDE



Python was created to be an easy scripting language to learn, and it is very forgiving with syntax. It is probably one of the best languages that we can recommend to start scripting in because it is so forgiving, unlike Bash or Windows batch scripting. Python also doesn't lack in features. Like VBScript and PowerShell, Python is extensible and can use external libraries. One disadvantage to using Python is that it's not been widely adopted in enterprise environments. This means you will not be able to build login scripts with it, because the interpreter does not have direct ties into the operating system. You'll often find that your flexibility is limited when you try to integrate a Python script into an enterprise process.

Python scripts normally end with the `.py` extension so that the end user can identify the scripting language used within the file. The `.py` extension also allows the script to be launched on Windows operating systems. With Linux/UNIX and macOS systems, however, the hashbang defines the interpreter to use and generally looks something like `#!/usr/bin/python3`.

This next script was written in Python and follows the same functionality as the previous examples in this chapter. As you can see, the program's readability is similar to the earlier examples and the syntax differs slightly.

```
#FlipFlop Script
FLIPFLOP = 0
while FLIPFLOP < 2:
    if FLIPFLOP == 0:
        print("Zero")
        FLIPFLOP = 1
    else:
        print("One")
        FLIPFLOP = 0
```

JavaScript

JavaScript and the Java programming language have several similarities—for example, they are both object-oriented languages, and their syntax is similar. However, the similarities end there, because Java is a programming language and JavaScript is a scripting language, with the difference being that Java is compiled whereas JavaScript is a scripting language.

JavaScript is mainly interpreted in web browsers to allow for interactive web pages. JavaScript is one of the three core web technologies; the other two are *Hypertext Markup Language (HTML)* and *Cascading Style Sheets (CSS)*. JavaScript can also be adapted to run outside the web browser with a runtime called `Node.js`. Both web browser JavaScript and `Node.js` scripts end in the `.js` extension, which identifies the contents of the file as JavaScript code.

JavaScript can be edited with any text-based editor. Microsoft Visual Studio Code does an excellent job of formatting for Windows-based editing of JavaScript. Brackets is another popular editor for JavaScript and runs on a variety of platforms.

The following is an example of JavaScript code. The example is similar to the previous examples in relation to functionality. This particular script was coded to run in `Node.js`, since JavaScript normally outputs to HTML. The structure is similar to the previous examples. As always, every language differs slightly in syntax.

```
//FlipFlop Script
FLIPFLOP = 0;
while (FLIPFLOP < 2) {
  if (FLIPFLOP == 0) {
    console.log ('Zero');
    FLIPFLOP = 1;
  } else {
    console.log ('One');
    FLIPFLOP = 0;
  }
}
```



For the 220-1202 exam, CompTIA does not expect you to write scripts. However, you should be able to identify scripts and know their associated extensions.

Scripting Use Cases

Scripting is a fairly new objective for the CompTIA A+ exam, and it may look overwhelming to you. You are not expected to have mastered the skills of writing scripts, such as those included in this chapter. They have been included so that you can visualize the various languages and understand the syntax. You will be required to read and understand what a script is doing, as well as identify the various elements of a script, such as variables, branch logic, basic loops, and the type of variables.

In addition to basic comprehension of scripts, you will need to know the various use cases where you may find yourself scripting something together. The basic rule of thumb should always apply: If the task is repetitious and needs to be completed several times, then a script should be developed. For example, if you needed to create 20 users, then a script is your best choice. Time versus reward should be calculated based on the time it takes to develop the script compared to the time it will take to complete the job. One consideration is how often you must complete the task; if you have three users you need to create every week, then developing a script is a good investment of time.

In this section, we'll examine several different scenarios in which you might find yourself developing scripts to complete.

Automation

The most compelling situation for the use of a script is one that requires some form of automation for a task. Scripts are perfect for *basic automation* tasks, such as creating users, adding users to groups, or even more intricate and sophisticated tasks. When a task is automated with a script, it guarantees that the task will flow the same every time it is executed.

There are two types of scripts that you will most likely create: scripts for automating your own tasks and scripts that automate tasks for others. During your career you will most likely find tasks that you have to do over and over. These tasks should be automated as much as possible, and each repetitive task should have its own script. As a best practice, you should create a folder that contains all the scripts you use on a daily basis. This way, you always know where they are, and when you move to a new computer, you can simply copy them over. Obviously, when you create scripts for others, you won't use them on a daily basis. However, these scripts should also be grouped together in a common folder, since you will probably reuse a part of one script or the entire script for another user.

The following is an example of a task that should be automated. It assumes that you have the Remote Server Administration Tools (RSAT) installed on your system, which includes the `dsquery` and `dsget` commands. The command line will query Active Directory for users with `test` in their name. The output from the initial command of `dsquery user -name *test*` accomplished that task and outputs the distinguished name (DN) of the user. We then pipe that output to the `dsget user -samid -ln -fn -email` command line and that retrieves the username, first name, last name, and email address of each user.



Remote Server Administration Tools (RSAT) is a feature that can be added to the Windows operating system to allow you to administer Windows Server roles such as Active Directory. These tools are not an objective of the CompTIA A+ exam but are useful to know. To learn more about RSAT, visit:

<https://learn.microsoft.com/en-us/troubleshoot/windows-server/system-management-components/remote-server-administration-tools>

```
C:\sys>dsquery user -name *test* | dsget user -samid -ln -fn -email
samid          fn          ln          email
testuser1      user1       test        test.user1@wiley.com
testuser2      user2       test        test.user2@wiley.com
testuser3      user3       test        test.user3@wiley.com
C:\sys>
```

This is a handy command line, but typing it in every time you need an answer like this is tedious. Luckily, we can write a simple script where we just need to change one part of the script. Where we had `test`, we simply replace it with `%1` to capture the first argument from the Windows batch script. Then we just start Notepad and copy the line in and save it as `lookup.cmd`.

```
dsquery user -name *%1* | dsget user -samid -ln -fn -email
```

Then when we run it, we get the following output:

```
C:\sys>lookup test
C:\sys>dsquery user -name *test* | dsget user -samid -ln -fn -email
```

```

samid      fn      ln      email
testuser1  user1    test    test.user1@wiley.com
testuser2  user2    test    test.user2@wiley.com
testuser3  user3    test    test.user3@wiley.com
C:\sys>

```

The script works just like before, but each command is echoed to the console. This is easily fixed by editing the file `lookup.cmd` with the following additional lines:

```

@echo off
dsquery user -name *%1* | dsget user -samid -ln -fn -email

```

Now when the file is executed, you won't see the echo of the command actually executed, but only the output of the command, as shown here:

```

C:\sys>lookup test
samid      fn      ln      email
testuser1  user1    test    test.user1@wiley.com
testuser2  user2    test    test.user2@wiley.com
testuser3  user3    test    test.user3@wiley.com
C:\sys>

```

We can refine the script further to add some branch logic, so if the user doesn't supply an argument, it explains the argument required:

```

@echo off
if {%1}=={} (
echo.
echo You must supply the following.
echo ex. %0 {name of user}
echo.
goto :END
)
dsquery user -name *%1* | dsget user -samid -ln -fn -tel -email
:END

```

The process of refining the script is called the development process. The first script you build will have common parts that you will reuse, such as error handling when no argument is supplied. In this example we used a Windows batch script, but you should use whichever scripting language(s) you are comfortable with. There is no right way to develop a script, and the purpose of this example is to show you the thought process involved in writing a script.

Restarting Machines

Another common task might be to restart multiple machines. You can solve this problem in several different ways, and it's all about the end goal of the task. If this is a one-time task, then you could write a simple script, such as the following:

```
shutdown /s /m \\computer1 /t 0
shutdown /s /m \\computer2 /t 0
shutdown /s /m \\computer3 /t 0
shutdown /s /m \\computer4 /t 0
```

If this is a recurring task, then you could develop a script that reads a text file into a variable and restarts each computer. It sounds complex, but it really isn't when you break down the task into smaller pieces. For example, say you need to read a file of computer names, and then walk through each line in the files and restart the computer. For this example, let's switch to PowerShell, since it allows reading of files and parsing them. Start by creating a text file named **list.txt** with four computer names in it: computer1 through computer4. Now we'll test the script by saving it as `restartcomps.ps1`.

```
$list = Get-Content .\list.txt
ForEach($line in $list) {
    Echo $line
}
```

The script will read the contents of the `list.txt` file into the variable `$list`. Then the script will load the variable `$line` for each entry in the `$list` collection. When we run the script, we get the following output:

```
PS C:\sys> Set-ExecutionPolicy Unrestricted
Execution Policy Change
The execution policy helps protect you from scripts that you do not trust.
Changing the execution policy might expose you to the security risks described
in the about_Execution_Policies help topic at
https://go.microsoft.com/fwlink/?LinkID=135170. Do you want to change the
execution policy?
[Y] Yes Yes to All [N] No [L] No to All [S] Suspend [?] Help (default
is "N"): y
PS C:\sys> ./restartcomps.ps1
computer1
computer2
computer3
computer4
PS C:\sys>
```

Now that we have the first two parts of the script working, we just need to add the functional piece. While we are refining the script, let's also remove the variable and read the lines directly:

```
foreach($line in $(Get-Content .\list.txt)) {
    Write-Output $line
    Restart-Computer -Computer $line
}
```

When the script is run, the output will remain the same, but you'll also be restarting the computers that are written as output to the console. We can make the script fancier by allowing arguments for the file to read, but we think you have the idea behind the script development process. It is a work in progress, and you always have some outstanding code to automate your job.

Remapping Network Drives

Remapping network drives can be done in a multitude of ways with VBScript, PowerShell, Window batch scripts, or some other favorite language. However, Windows batch script and PowerShell are the most common. In the following, we see a script that maps a few drives using a Windows batch script. When we map a drive, we are mounting a remote filesystem through to a drive letter. For example, we can mount the remote filesystem of \\server1\files to a local drive letter of m:

```
net use m: \\server1\files
net use n: \\server2\files
net use o: \\server3\files
```

The same script can also be developed in PowerShell:

```
New-PSDrive -Name "m" -PSProvider FileSystem -Root "\\server1\files"
New-PSDrive -Name "n" -PSProvider FileSystem -Root "\\server2\files"
New-PSDrive -Name "o" -PSProvider FileSystem -Root "\\server3\files"
```

These scripts don't need to be complex like the previous scripts. We just want to obtain a reproducible result every time the script runs.

Installation of Applications

When scripting is combined with the installation of applications, you can perform a number of functions that are not possible on their own. For example, you can write a script that installs the prerequisites for an installation and then only succeeds if the subsequent installations are successful. A lot of these scenarios are going to be custom to your specific needs and environment. The following is an example of a PowerShell command that will install an application called App.msi:

```
Invoke-CimMethod -ClassName Win32_Product -MethodName Install -Arguments
@{PackageLocation='\\server\installs\App.msi'}
```


A few assumptions are made with the PowerShell example. The first assumption is that you are calling an MSI installer. If you aren't, the code will not work, since every installer has its own methods for invoking an installation. The second assumption is that you are an administrator of the operating system. Writing a PowerShell script will not circumvent Windows security.

Creating a scripted installation of an application is rewarding, but it is also time consuming. You will most likely have to refine your script several times before it works as expected. This means that you will need to install the application many times to get it right. However, if you have an application that requires installation across a number of computers, you can easily reclaim the time spent on the script.

Automated Backups

Backups should be trusted to back up software that is engineered to expire media, rotate media, and generally backup and restore the data and systems that the organization depends on. This type of software is considered off-the-shelf backup software and scripts are not expected to replace this software. However, by using scripts you can automate pieces of the backup process to make the process much more reliable.

A common example of automating backups with scripts is the backup of SQL databases. You can use off-the-shelf backup software to back up the database. However, one common problem is that the agent installed on the SQL server that facilitates the backup will take a snapshot of the database and back it up in whole. This might seem like what we are trying to achieve as an end goal—except when you try to restore it, you'll quickly find that you need to restore the database in whole, even though you only need one table of records.

If you preprocess the database with a script, you can export it to a file. This file can then be backed up and restored to any SQL server, even down to the record level. The following is a maintenance script used with Microsoft SQL to back up the database to a file. The script is written in the SQL scripting language.

```
USE TestRecords
GO
BACKUP DATABASE [TestRecords]
TO DISK = N'D:\DBBackups\TestRecords.bak'
WITH CHECKSUM;
```

You can automate many other types of backups, such as custom application data, Registry settings, email data, and any other type of data you can address with a script. However, one of the most common backups you will automate with scripts are SQL database backups.

Gathering of Information/Data

Although you can use applications such as Device Manager to gather information about a computer, using the GUI is not scalable when you need to gather information from a large group of computers. This is where scripts come in handy to gather information and even export the information to a file for you. In the following examples, we will look at some

simple PowerShell commands for gathering information and exporting the information. However, you can use any scripting language you like and Windows batch scripting is very commonly used as well.

```
PS C:\sys> Get-Service
```

Status	Name	DisplayName
Running	agent_ovpnconnect	OpenVPN Agent agent_ovpnconnect
Stopped	AJRouter	AllJoyn Router Service
Stopped	ALG	Application Layer Gateway Service
Stopped	AppIDSvc	Application Identity
Stopped	Appinfo	Application Information

[Output Cut]

The `Get-Service` cmdlet will show you all the services running on the operating system. It will output a long list and will display each service's Status, Name, and DisplayName. By piping the output to the `Export-Csv` cmdlet, we can output a lot more detail and send it directly to a comma-separated values (CSV) file. An example of this is in the following command string:

```
PS C:\sys> Get-Service | Export-Csv .\Services.csv
```

The `Get-Service` cmdlet isn't the only command you can use to gather data; there are various other commands that let you gather data and export it. Obviously, you can chain these commands together in a script and gather a large amount of data. You can then use scripts to mine the data for specific information, such as a service state or free space.

Initiating Updates

The Windows platform has been automatically patching itself since the release of Microsoft Update (now called Windows Update) and Windows XP. The feature of Microsoft Update and the Windows platform have evolved since the early release. It has turned into a robust feature to keep Windows up to date. However, if you need an immediate result to patch a security hole, then scripting a solution is the best remedy.

There are a number of ways to get patches to immediately install on Windows. The approach you choose depends on your patch management solution, such as Windows Update, Windows Server Update Services (WSUS), Microsoft Endpoint Configuration Manager (MECM), or a third-party patch management solution. For the remainder of this section, we will use the Windows Update solutions for examples.

If you want to script Windows Updates to initiate patching via Windows batch scripting, then the utility of choice is `wuauclt.exe`. The command can be directed to detect patches with the `/detectnow` argument. However, don't expect anything elaborate; the utility will not notify you that it is doing anything. If you want to watch the progress, then you'll need to keep an eye on the log file, `C:\Windows\WindowsUpdate.log`.

An alternative is to use PowerShell and a PSGallery module call `PSWindowsUpdate`. This module allows you to fully automate the patch management and obtain a great level

of detail. For example, after installing the module you can execute the command `Get-WindowsUpdate` to obtain the pending list of available updates, as shown here:

```
PS C:\sys> Install-Module -Name PSWindowsUpdate
NuGet provider is required to continue
PowerShellGet requires NuGet provider version '2.8.5.201' or newer to interact
with NuGet-based repositories. The NuGet provider must be available in
'C:\Program Files\PackageManagement\
ProviderAssemblies' or 'C:\Users\bohack\AppData\Local\PackageManagement\
ProviderAssemblies'. You can also install the NuGet provider by running
'Install-PackageProvider -Name NuGet
-MinimumVersion 2.8.5.201
-Force'.
Do you want PowerShellGet to install
and import the NuGet provider now?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): y
Untrusted repository
You are installing the modules from an untrusted repository. If you trust this
repository, change its InstallationPolicy value by running the Set-
PSRepository cmdlet. Are you sure
you want to install the modules from 'PSGallery'?
[Y] Yes Yes to All [N] No [L] No to All [S] Suspend [?] Help (default
is "N"): y
PS C:\sys> Get-WindowsUpdate
ComputerName Status      KB          Size Title
-----
CERES          ----- KB5007406    80MB 2021-11 Cumulative Update Preview for
.NET Framework 3.5, 4.7.2 and 4.8 for...
CERES          ----- KB2267602    2GB Security Intelligence Update for
Microsoft Defender Antivirus - KB2267602 (...
PS C:\sys>
```

You can then execute the command `Get-WUInstall -AcceptAll -AutoReboot` and the operating system will begin to install the updates and automatically reboot. These are just a few ways you can script the installation of Windows Updates.

Other Scripting Considerations

Along with the knowledge of scripting comes great power and even greater responsibility. You should consider several key points before scripting a solution and during the development of the script. In the following section we will cover some of these key considerations.

Of course, there may be factors outside of these CompTIA objectives. The one consideration that has resonated throughout the previous section is the decision between investing time to develop a script or just completing the task. This is something that you will need to take into account before you even begin scripting.

Introducing Security

A common pitfall with scripting is inadvertently introducing a security issue. Security issues come in all different forms when scripting. The most common security issue is the embedding of security credentials in scripts. Regardless of how secure you think the script will be, it's a bad habit and should be avoided at all costs.

That being said, there are instances where embedding a password cannot be avoided. In these situations, you should use mechanisms that are supported in the scripting language to act as a digital locker for your password or methods that encrypt the password. One such mechanism is the `ConvertTo-SecureString` cmdlet. It means more lines of code, more time, and sometimes more aggravation in getting it to work, but the benefit is a secure system.

Another potential problem is the unintentionally introduction of malware with scripting. This usually happens when you need additional functionality, such as email capability or changing a system setting. Windows batch scripting has its limitations, and it is easy to use a third-party tool to get the last bit of functionality out of the script. However, that third-party tool may have malicious intent, and you could introduce that malicious code to the system in which the script is executed. PowerShell has its own potential malware risks; when you add an untrusted module, the same could happen.

The best way to avoid these pitfalls is to not shortcut the solution. Do your homework and avoid embedding passwords or using untrusted modules or third-party utilities in your scripts. It means more work, but this should be factored into your cost-benefit analysis when taking on developing a script.

Changing System Settings

Scripts are generally created to change various user and system settings. However, sometimes you can inadvertently change a system setting that is not meant to be changed. If this happens, it can adversely affect the system it is executed on. Therefore, it is always best to test the script on a sacrificial computer. A virtual machine is the best option, since it can be reverted to a prior snapshot if something unexpected happens.

An example of a simple mistake is using the `setx` command incorrectly. The `setx` command modifies the system environment variables and is used by Windows batch scripting. By using the statement `setx path "$path;c:\sys" /m`, you could cause irreversible problems. This statement will set the path variable to `$path;c:\sys`, because the incorrect syntax was used to address the variable of path. The correct statement is actually `setx path "%path%;c:\sys" /m`. The statement now sets the path variable to the current path variable contents (`%path%`) and the addition of `c:\sys`.

The mistake can be as simple as mixing up syntax from one scripting language with syntax from another scripting language. However, some commands are not as forgiving, such

as the previous example. When using PowerShell the mistake can be very subtle, since many commands cover both user and system settings. Always test your script before you execute it on a live system. Otherwise, you might be reinstalling an operating system in addition to completing your scripting task.

System Crashes

Scripting allows you to automate processes, which is the main reason we are developing the script in the first place. However, automation can sometimes create problems that we can't foresee during the development of the script. As an example, you can quickly gobble up all the usable RAM resources with the following Windows batch script. The script will launch instances of Notepad until the operating system runs out of memory.

```
:loop  
start notepad.exe  
goto :loop
```

Although this script is obvious in its intent, it is an extreme example of automation that if left unchecked will crash the system. Windows batch scripting is not the only scripting language where things can go awry; you can do the same with PowerShell. This example will create an HTML-formatted file of the directory structure:

```
Get-ChildItem c:\ -Recurse | ConvertTo-Html | Out-File -FilePath .\output.html
```

The problem with this example is the sheer size of the resulting file. This statement will recursively list all the files from the C: drive down. When the file is launched in the web browser, the browser will quickly run out of memory trying to display the large file. More elaborate scripts may automatically open the web browser and immediately crash it.

To prevent similar problems from happening in your environment, you should test and monitor your scripts. By testing your script solution for errors or conditions that can run the system out of resources unintentionally, you identify and correct problems that would otherwise cripple the system. Monitoring should be performed after the script is in place, such as a script that is started via the Task Scheduler. You want to identify any spikes in resource usage after the solution is in place.



Exam Objective 4.9, "Given a scenario, use remote access technologies," requires you to know the various methods and tools used for remote access. This includes third-party tools like screen-sharing, videoconferencing, file transfer, and desktop management software.

Remote Access


Remote access technologies have been around for as long as networks have been in existence. As a technician, you can't always be in all the places you are needed.

Remote access technologies allow you to do just that. You can create a remote connection to a distant computer on the network and change its configuration. In this section, we will discuss several different types of remote access technologies that are covered on the CompTIA 220-1202 exam, as well as their security implications for a network.

Remote Desktop Protocol

Remote Desktop Protocol (RDP) is used exclusively with Microsoft operating systems. An RDP client—the *Remote Desktop Connection client*, also known as the `mstsc.exe` utility—is built into the Microsoft operating system, as shown in Figure 8.6. The Remote Desktop Connection client can provide remote access as though you were sitting in front of the keyboard, monitor, and mouse. The Remote Desktop Connection client and RDP can transport several other resources, such as remote audio, printers, the clipboard, local disks, and video capture devices, as well as any other Plug and Play (PnP) devices.

FIGURE 8.6 The Remote Desktop Connection client



RDP communicates over TCP port 3389 to deliver the remote screen and connect the local mouse and keyboard for the RDP session. RDP uses Transport Layer Security (TLS) encryption by default, and it provides 128-bit encryption. Microsoft allows one remote user connection or a local connection on desktop operating systems via RDP, but not both

at the same time. On server operating systems, Microsoft allows two administrative connections, which can be a combination of local or remote access but cannot exceed two connections.

Microsoft also uses RDP to deliver user desktops and applications via terminal services. When RDP is used in this fashion, a centralized gateway brokers the connections to each RDP client desktop session. Terminal services require terminal services licensing for either each user connecting or each desktop served. RDP can also be used to deliver applications to end users using Microsoft RemoteApp on terminal services. When RemoteApp is used, the server still requires a terminal services license. However, just the application is delivered to the user rather than the entire desktop.

Virtual Private Network (VPN)


A virtual private network (VPN) extends your company's internal network across the Internet or other unsecured public networks. This remote access technology allows clients and branch networks to be connected securely and privately with the company's network. There are several different ways a VPN can be used in your network architecture, and we will cover them in the following sections. A VPN achieves this private connection across a public network by creating a secure tunnel from end to end through the process of encryption and encapsulation. The encryption protocols used vary, and we will cover them as well. Since a tunnel is created from end to end, your local host becomes part of the company's internal network along with an IP address that matches the company's internal network. We don't have to be bound to only TCP/IP across a VPN, since this technology can encapsulate any protocol and carry it through the tunnel.

Site-to-Site VPN

Over the past 10 to 15 years, using high-bandwidth connections to the Internet has become cheaper than purchasing dedicated leased lines. So, companies have opted to install Internet connections at branch offices for Internet usage. These lines can serve a dual purpose: connecting users to the Internet and connecting branch offices to the main office. However, the Internet is a public network and unsecured, but site-to-site VPN connections can fix that. Companies with multiple locations have reaped the benefits of creating VPN tunnels from site to site over the Internet by ditching their leased lines, installing VPN concentrators at each location, and creating VPN tunnels. Site-to-site VPN is also much more scalable than leased lines because locations only need a connection to the Internet and a VPN concentrator to be tied together. Figure 8.7 details two locations tied together with a VPN tunnel. The magic happens all in the VPN concentrator. Since VPN concentrators also have a routing function, when a tunnel is established, a route entry is created in the VPN concentrator for the remote network. When traffic is destined for the branch office with a destination network of 10.2.0.0/16, the router encrypts and encapsulates the information as data and sends it to the other side of the tunnel over the Internet. This is similar to a host-to-site VPN, the difference being the routing is performed in the VPN concentrator. When the packet is received on the other side of the tunnel, the VPN concentrator decapsulates the data, decrypts the packet, and sends the packet to its destination inside the branch network.

It is common to find that the appliance performing VPN is also the firewall and router. Firewalls today are sold with VPN software built in and licensed accordingly.


FIGURE 8.7 A typical site-to-site VPN



Client-to-Site VPN

Client-to-site VPN connectivity is a remote access strategy for mobile access. It can be used for telecommuters, salespeople, partners, and administrative access to the internal network resources. The key concept is that VPN access is granted on an individual or a group basis for the mobile users. Using the example in Figure 8.8, you can allow salespeople to connect to the corporate network so they can update sales figures or process orders. This can all be done securely over the Internet while the users are mobile and have access to a network connection.


FIGURE 8.8 A typical host-to-site VPN



When a client computer establishes a VPN connection, it becomes part of the internal corporate network. This happens by assignment of an IP address from the internal corporate network. In Figure 8.9, you can see a mobile device such as a laptop with a VPN client

installed in the operating system. When the connection is established with the VPN concentrator over the Internet, a pseudo network adapter is created by the VPN client. In this example, the pseudo network adapter is assigned an internal IP address of 10.2.2.8/16 from the VPN concentrator. The laptop also has its own IP address of 192.168.1.3/24, which it uses to access the Internet. A routing table entry is created in the operating system for the 10.2.0.0/16 network and through the pseudo network adapter. When traffic is generated for the corporate network, it is sent to the pseudo adapter, where it is encrypted and then sent to the physical NIC and sent through the Internet to the VPN concentrator as data. When it arrives at the VPN concentrator, the IP header is stripped from the packet, the data is decrypted, and it is sent to its internal corporate network resource.

FIGURE 8.9 Client-to-site VPN connection



Clientless VPN

There are many different VPN solutions on the market. Each one of them traditionally requires the installation of a VPN client. However, there are a growing number of products that do not require the installation of a client; these products are called clientless VPN solutions. The VPN client is the web browser on the mobile device requiring connectivity back to the corporate network. The VPN appliance acts as a reverse proxy to the various resources internal to the organization.

Virtual Network Computing

Virtual Network Computing (VNC) is a remote-control tool for the sharing of desktops. The VNC client normally operates on TCP port 5900. VNC is similar to Microsoft RDP, with the exception that VNC is an open source protocol and typically allows only one console session per operating system. It supports encryption via plug-ins, but it is not encrypted by default.

VNC operates in a client-server model. The server allows for the remote control of the host on which it is installed. It is normally configured with a simple shared password, but it can also be configured with Windows groups. Several different clients can be used, such as RealVNC, TightVNC, and many others, but they all perform similarly.




Many different VNC clients are available for download. To learn more, visit www.realvnc.com and www.tightvnc.com. A simple web search for “VNC clients” will produce a staggering number of results. Each client has its own unique feature set.

Secure Shell

Secure Shell (SSH) is commonly used for remote access via a text console for Linux and UNIX operating systems. The SSH protocol encrypts all communications between the SSH client and the SSH server using TCP port 22. The SSH server is also known as the SSH daemon. SSH uses public-private key pair cryptography to provide authentication between the SSH client and server. SSH can also use a key pair to authenticate users connecting to the SSH server for the session, or a simple username and password can be provided.

It is important to understand that both the user and their host computer are authenticated when the user attempts to connect to the server. During the initial connection between the user’s host computer and the SSH server, the encryption protocol is negotiated for the user’s login, and the cryptography keys are verified. PuTTY, shown in Figure 8.10, is a common SSH client that is free to download and use. PuTTY provides various methods of connecting to a remote device, including Telnet and SSH.

FIGURE 8.10 PuTTY SSH Client



Microsoft included the OpenSSH client by default in Windows 10 version 1803 and later. The OpenSSH client is located in the `C:\WINDOWS\SYSTEM32\OpenSSH` folder.

Beyond logging into a Linux or UNIX server for remote access, SSH can provide remote access for applications. Through the use of SSH port forwarding, the application can be directed across the SSH connection to the far end. This allows applications to tunnel through

the SSH session. It also encrypts application traffic from the client to the server, because it is carried over the SSH encrypted tunnel. SSH can behave similarly to a VPN connection, but it is more complex to set up.

Remote Monitoring and Management (RMM)

In the early days of your organization, it may have been simple to monitor and manage all the various systems from one location. However, as your organization's footprint grew across different sites and many employees now work from home (WFH), it is more difficult to monitor and manage the various systems. Systems need patching, must be monitored for disk space, and have hardware and applications installed—and these are just a few of the tasks.

This is where a *Remote Monitoring and Management (RMM)* solution can help IT across your enterprise or multiple enterprises and give you a holistic view of your enterprise. There are several different RMM solutions on the market today. Among the most popular solutions are managed service providers (MSPs) that manage your enterprise for a contracted price. These service providers ultimately use RMM software to monitor and maintain the enterprise. The MSP will require your organization to install an agent that is configured to report back to the MSP's RMM software.

You can also purchase a cloud-based or on-premises solution for maintaining your organization with your own IT department. Every vendor of RMM solutions has their own variation of features, which makes up the product's secret sauce. These solutions also require the installation of an agent that reports back to the RMM solutions. Regardless of which product you choose, there are two main features to RMM: the remote monitoring feature and the management feature.

Remote Monitoring The remote monitoring feature of an RMM system can monitor a number of different components, such as security, hardware, applications, and even activity on the operating system. These are just a few of the components of the remote monitoring feature found in RMM systems. The list grows depending on the RMM vendor. The most common monitoring is the security of the various systems across your enterprise, such as patch levels, antimalware status, and exploits. Another monitored component is the hardware and applications installed across the enterprise. Monitoring the hardware can identify your assets, as well as identify when upgrades are needed. Application monitoring can identify problems with a specific application or your vulnerability in the event the application needs to be patched. These are just a few monitored components—the list grows with every release of new RMM software by vendors.

Reporting is a major component of the remote monitoring capabilities of RMM systems. The reporting can be active or passive for most systems. In an active reporting system, the RMM software will compile a report periodically and alert you when a major change is discovered. As an example, if over 30 percent of your computers are vulnerable to a new exploit, the system can be configured to alert you. You may also set up a similar threshold alert for disk space. The passive reports can be run and give you

an overall picture of your network and are typically in the form of a drill-down report. A drill-down report allows you to view the overall health and drill down to specific areas of interest to a specific detail.

Remote Management The remote management feature of RMM solutions can also be quite diverse in the type of management they provide. Many RMM vendors offer the management of computers via a form of remote desktop control. This remote desktop control allows you to intervene when a user has a problem. RMM solutions also offer integrated patch management solutions to patch applications and the operating system. In addition to user intervention and security, you may find remote application installation, disk cleanup, and remote antimalware scan features, and these features just scrape the surface of RMM capabilities. Every vendor of RMM systems has their own unique management capabilities based on their product's focus, such as security, asset management, or user productivity.

Simple Protocol for Independent Computing Environments (SPICE)

The Simple Protocol for Independent Computing Environments (SPICE) is an open remote computing solution that provides high-quality and rich user experience for remote access to virtual desktops. The solution was originally developed by Qumranet in 2007 and later purchased by Red Hat. The original goal of SPICE was to enhance the interaction between users and virtual machines by offering a seamless and efficient protocol. SPICE provides a framework that enables users to connect to their virtual desktops securely and efficiently from any location.

SPICE is integrated with various virtualization platforms, such as Proxmox VE. This integration allows users to connect to the virtual desktop infrastructure (VDI) environment Proxmox VE hosts. The user can use keyboards, mice, audio, and video, with the virtual desktop as if these devices were connected to their local device connected to the virtual desktop. The experience is the same if the user is connected on a LAN or WAN connection.

SPICE operates by establishing a connection between the client device and the virtual desktop hosted on a server. The connection process involves several components, including the SPICE server, SPICE client, and SPICE protocol. The SPICE server resides on the host machine (usually a virtualization server) and handles the virtual machine's display and input/output operations. The SPICE client, installed on the user's device, is responsible for establishing the connection to the SPICE server and rendering the virtual desktop's display locally.

One of the features of SPICE is its robust security framework. The protocol supports multiple data channels, each of which can operate in either clear-text or TLS modes. This flexibility allows administrators to balance security and performance based on their specific needs. Additionally, SPICE offers a range of authentication schemes, including Kerberos.

Windows Remote Management (WinRM)

Windows Remote Management (WinRM) is a powerful tool used in Windows operating systems to facilitate remote management and automation of tasks. It allows administrators to securely execute commands, manage configurations, and retrieve data from remote computers. The initial WinRM protocol version 1.1 was originally shipped with Windows Vista and Windows Server 2008. Currently the WinRM protocol has matured to version 3.0 and is an integral part of the current Windows operating systems. WinRM is built on the WS-Management protocol and uses the Simple Object Access Protocol (SOAP). SOAP is a protocol used for exchanging information that is formatted in an XML format. WinRM also incorporates security protocols TLS, Kerberos, and NTLM (NT LAN Manager) to secure the WS-Management protocol. The WinRM protocol operates over standard HTTP port 80 and HTTPS port 443.

One of the most powerful features of WinRM is its ability to execute commands and scripts remotely. Administrators can use WinRM to run PowerShell scripts on remote systems, enabling them to perform a wide range of administrative tasks without needing physical access to the machines. This capability is particularly useful for automating repetitive tasks, such as software deployment, system updates, and configuration changes. By leveraging WinRM, administrators can write scripts that perform complex operations across multiple systems simultaneously, significantly reducing the time and effort required for system management.

WinRM also plays a crucial role in monitoring and troubleshooting remote systems. Administrators can use WinRM to gather performance data, check system logs, and diagnose issues without interrupting the end user's activities. This remote diagnostic capability helps in quickly identifying and resolving problems, minimizing downtime and maintaining system reliability. For example, an administrator can use WinRM to remotely check the event logs of a server to identify the cause of a system crash or performance issue.

Remote Access Tools


Since the release of Windows XP, Microsoft has included various tools to allow remote assistance to the Windows operating system. In addition to Microsoft's proprietary remote assistance tools, many vendors have entered the market. As a result, there are a number of third-party remote assistance tools freely available and out on the market, with varying costs. Let's explore some of the built-in capabilities of the Windows product and some of the features of third-party products.

Quick Assist

Microsoft includes a remote access tool in Windows 11 for soliciting help called *Quick Assist*, shown in Figure 8.11. Quick Assist operates similarly to third-party services such as Splashtop, GoToMeeting, and Join.me. These tools allow for remote desktop sharing and assistance for a user through screen sharing. Unfortunately, Quick Assist does not support

file transfers. Quick Assist does work behind routers and firewalls, unlike prior built-in applications. Quick Assist is becoming a tool of choice for many helpdesk tasks. Quick Assist even has its own shortcuts for easy access by the user.

FIGURE 8.11 Windows Quick Assist tool



You can launch Quick Assist on Windows 11 by clicking the Start menu and typing **Quick Assist**. Once the Quick Assist tool is started on your computer (the assistant), choose Give Assistance. You will then be prompted to log in with a Microsoft account. Once you are signed in, you will be presented with a six-digit number, which will be valid for 10 minutes. On the end user's (person being helped) computer, the end user will launch Quick Assist and select Get Assistance. The end user will need to obtain your six-digit code and enter it into the dialog box when prompted. Once the number is entered on the end user's side, both computers will connect, and you will be asked what type of control you want to provide for the end user: View Only or Full Control. Of course, the end user has to allow the control, but you can provide remote assistance without installing any third-party software.

Quick Assist offers chat functionality to the assistant and the end user in a chat window. You can launch this window by clicking Toggle Instruction Channel on the Quick Assist toolbar. You can also launch Task Manager automatically in later versions of the tool. By clicking the Task Manager item on the Quick Assist toolbar, you can remotely launch Task Manager on the end user's operating system. An annotation tool allows you to draw on the end user's screen to guide them. Select this tool by clicking the Annotate option on the Quick Assist toolbar. However, Quick Assist does not have a file transfer utility. Any file transfers

must be done with a file sharing service, where the user downloads the file from the Internet link you provide.

Third-Party Tools

Third-party tools such as Splashtop, GoToMeeting, and Join.me have their own unique features. Each tool fits into a category of screen sharing, videoconferencing, file transfer, or desktop management. When choosing a third-party tool, you should evaluate the main requirements and the category the software excels at. Let's examine the various categories to help you better identify the best software for your needs.

Screen Sharing Both the MSRA tool and Quick Assist fit into the category of *screen-sharing* software. Screen-sharing software will share your screen with another user, and it is most often used to troubleshoot problems. Two popular screen-sharing software utilities are Splashtop and TeamViewer. Both allow you to share your screen with another person, similar to the MSRA tool and Quick Assist. However, these tools have other built-in features that the MSRA tool and Quick Assist do not, such as file sharing and screen blanking.

Videoconferencing The *videoconferencing* category includes any screen-sharing software that allows more than one person to screen-share. The software also has some additional functionality, such as calendar invite, call-in numbers, and attendee management. Microsoft Teams, Webex, Zoom, and GoToMeeting are all examples of videoconferencing software. These tools are often used to present ideas and collaborate between groups of people.

File Transfer The ability to file transfer between two computers is not really a category of software; it is a feature of screen-sharing and/or videoconferencing software. *File transfer software* can, however, be a category of remote access software if you are solely looking for that feature. Popular file transfer software includes OneDrive, Dropbox, and Google Drive. These file transfer utilities allow the sharing of files between users and serve a specific purpose for the distribution of files bidirectionally.

Desktop Management The last category is an extension of some screen-sharing software packages such as Splashtop and TeamViewer. Both of these products can be used in an ad hoc fashion for screen sharing. They can also be installed on the operating system and configured to allow an inbound connection. When these tools are installed to await an incoming connection, they are considered to be *desktop management software* solutions. Many of these tools also allow integration with other management tools like Intune and Jamf. This integration gives you complete device management, and a lot of these tools will even show when a device is online and connectable.

Security Considerations

Each of the remote access technologies discussed in this chapter have security considerations. Before implementing a remote access technology, you should determine what type of data is

going to be exchanged and whether the level of encryption is sufficient. Telnet, which is not encrypted, might be fine if simple data is being transmitted, such as temperature or humidity readings. However, if passwords, configuration, or any type of sensitive information will be transmitted, then a more secure protocol, such as Secure Shell, should be used.

Beyond the data in transit and the method that provides the transit, there are other security considerations. One of the biggest concerns is any agent that awaits connection and is exposed to the Internet. A threat agent can exploit these software packages and compromise a host via these software packages. To combat this problem, keep your software package up to date, and if multifactor authentication is available, use it.

Videoconferencing packages should be secured so that a password is required to join a meeting. Setting a password will thwart conference bombing, also known as *Zoombombing*. This is the act of a threat agent guessing the meeting ID and joining an otherwise private conversation. This type of security concern isn't really your typical threat, but it is a liability to your organization if not avoided.



Zoombombing became an annoyance on the Internet during the COVID pandemic. Organizations were disbanded from their campuses and employees were sent home to work. Because Zoom conferencing software is a dominant software product for videoconferencing, many organizations used the software to continue doing business. With the higher usage and the relaxation of security by organizations, Internet trolls guessed meeting IDs and joined otherwise private conferences. The trolls mainly caused disruption to the ongoing conferences, but many exposed the participants to unsavory behavior.



Exam Objective 4.10, "Explain basic concepts related to artificial intelligence (AI)," requires you to know AI concepts. You should know the basic components of AI systems, how they are trained, application, policies, limitations, and private versus public datasets.

Artificial Intelligence

The goal of artificial intelligence (AI) is to allow computers to simulate or imitate human intelligence. Although there are certain parallels between AI and human intelligence, their nature and capacity are essentially distinct. Consciousness, self-awareness, and emotional experience are hallmarks of human intellect, which is a biological evolution. Human intelligence enables people to think abstractly, be creative, adapt to new circumstances, and learn from their experiences. Human culture and social interactions are closely linked to our intelligence; this type of intelligence is sometimes referred to as emotional intelligence.

For the majority of us, it allows us to comprehend and react to intricate social signs and emotional subtleties.

AI is a product of human ingenuity and is intended to emulate specific cognitive processes like learning, thinking, and problem-solving. AI systems are excellent at processing large volumes of data fast and precisely, especially those driven by machine learning and deep learning. In some areas, they are more accurate and efficient than humans at recognizing patterns, making predictions, and carrying out tasks. AI can, for instance, create text that appears human, streamline supply networks, and identify ailments by analyzing medical photos.

The subject of AI has become a hot topic as it has become integral in our lives and complementing human intelligence. Math teachers often said, “You must learn math because you are not always going to have a calculator in your pocket.” Today, we not only have calculators as one of the many mobile apps, but we also have advanced AI that can help us solve problems. In this section, we will explore the components of AI, the characteristics of data AI uses, as well as the use of AI and considerations.

Evolution

Believe it or not, AI is not a new technology. In fact, the idea of neural networks and machine learning dates all the way back to the 1950s. John McCarthy and Alan Turing are claimed to be the original pioneers who introduced the concept of AI. In Turing’s groundbreaking work “Computing Machinery and Intelligence,” he raised the question of whether machines were capable of thought. The phrase *artificial intelligence* was first used by McCarthy, who also arranged the Dartmouth Conference in 1956. The Dartmouth Conference is regarded as the beginning of AI as a field of study.

In the preceding years, AI research helped fuel the subject and produced the first artificial neural network called the Perceptron. It was designed by psychologist Frank Rosenblatt in 1958. The Perceptron design was based on 1s and 0s and was a very simple proof of concept. The Perceptron would take input, weigh the input (which adjusted the input’s importance), use a bias to make a prediction, and create output based on the weighted sum of the bias and prediction. This research is rudimentary but similar to how AI works today. It introduced the concept of machine learning to produce the bias and prediction.

The computer boom of the 1960s and 1970s helped continue the progression of AI research. The first working form of AI on computers of this era was something called expert systems. Expert systems were nothing more than a computer program that mimics human experts. It did this by using data and basic programming that mimicked a thought pattern. As an example, some of these expert systems allowed for the diagnosis of patients and a treatment plan. These expert systems were far from the AI we know today, but they were an important stepping-stone to where we are now. The early decades eventually paved the way for development of foundational concepts and algorithms, such as symbolic reasoning and early neural networks. However, progress was slow due to limited computational power and data.

The continued work on expert systems and the growth of machine learning methods in the 1980s and 1990s helped the expansion of AI. However, AI was limited to rudimentary chatbots and less well-known systems—hardly considered AI. The real breakthrough came in the 21st century with the availability of large datasets and advancements in neural network architecture. This gave new breath to machine learning algorithms, because the data was there to feed these systems and thus better decision models could be made. Of course, machine learning algorithms got better as a result, but the large dataset was the catalyst. The best example of this generation of AI is IBM's Deep Blue. In 1997 Deep Blue went head-to-head with the reigning chess world champion Garry Kasparov. Deep Blue won the title of chess master after a six-game marathon! This was a milestone for AI since its inception in the 1950s.

The true revolution in AI came in the 2010s with the advent of deep learning, a subset of machine learning that utilizes multilayered neural networks. In the decade that followed, we have seen the advancement of several foundation models. These foundation models allow us to perform image recognition, word recognition, voice recognition, and ultimately natural language processing. Each of these models was a milestone in AI.

As early as 2006, Google introduced language translation, or simple text-to-text translation from one language to another. In 2016, Google introduced the Google Neural Machine Translation (GNMT) system, an advancement that allowed the language translation service to take images with words and translate them to any language. In 2022, Google launched Translation Hub, which provides document translations. In the following years, Google added more languages and larger datasets. This is just one example of the advancement of AI to date. There are many other examples, such as ChatGPT from OpenAI, Microsoft Copilot, Apple Intelligence, and Google Gemini. These are just a few popular AI models that use generative AI. AI just keeps getting better and better with each refinement of these milestones. Just imagine what the next 70 years will look like.

Components

The subject of AI, like many other complex topics, requires us to know everything before we can learn something. This section explains the various components of AI we discussed in the previous section. As you progress through this section, how AI works should become clearer.

Machine Learning (ML)

Machine learning (ML) is a broad field of AI that focuses on developing algorithms that allow computers to learn from and make predictions or decisions based on data. ML works by using algorithms to identify patterns and relationships within large datasets. These algorithms can then make predictions or decisions based on new data inputs. For example, if you were to memorize a series of numbers such as 11364, 90210, 89049 and were presented with the first three digits of 902, you could predict the last two, or 10. This is a type of ML called supervised learning.

There are three different types of learning displayed in learning models. The supervised learning model allows AI to predict or decide, based on historical data. In an

unsupervised learning model, the AI analyzes data to derive a new outcome. And the reinforcement learning model allows AI to properly sequence events and learn by reward. All of these models make up some of the human intelligence that we use daily as human beings. Neural networks, deep learning, and tokenization are other examples.

Supervised Learning

With supervised learning, the model is trained on a labeled dataset, meaning that each training example is paired with an output label. The model learns to map inputs to the correct outputs. Let's build off the prior example of 5 digit numbers being trained as sequences such as 90210. The label for these numbers would be ZIP codes and of course the output would be Beverly Hills, California. Another example is that it can be used for email spam detection. The model is trained on a dataset of emails labeled as "spam" or "not spam" and learns to classify new emails accordingly.

Unsupervised Learning

In a supervised model, the dataset contains labeled data, like spam or not spam. In an unsupervised model, the dataset does not contain labeled data. With unsupervised learning, the model is given data without explicit instructions as to what to do with it. This model is often used to identify hidden patterns or groupings within the data. As an example, if you owned a retail store and had a large volume of transactions, you could use this model to identify segments of customers. You could do this by grouping the customers into clusters of transactions and then labeling the output as frequent shopper, bargain hunter, or high-value customer based on their transactions.

Reinforcement Learning

This type of learning model is based on a sequence of decisions to a desirable outcome or rewards. The model learns to achieve a goal by maximizing cumulative rewards. This type of learning model is similar to a computer playing chess against a human. The sequence of decisions leading up to checkmate is analyzed for the cumulative reward. Another example is a self-learning robot navigating a maze, where each movement is calculated against the previous with the goal of solving the maze.

Neural Networks

Neural networks are a class of machine learning model that mimic the human brain's structure and function. Neural networks are made up of several layers of interconnected nodes that process and send data, similar to real neurons. Each neuron receives input and applies a transformation to the data using weights and biases; then it passes the output to the next layer. This process allows neural networks to learn complex patterns and relationships within a dataset and makes them highly effective for tasks such as image recognition, natural language processing, and game playing.

Neural networks play a crucial role by enabling machines to perform tasks that typically require human intelligence. There are several different types of neural networks depending on the task. For example, convolutional neural networks (CNNs) are specialized for image and video analysis. These neural networks allow AI systems to identify objects, faces, and

even emotions in visual data. Recurrent neural networks (RNNs) are good at processing sequential data, making them ideal for applications like language translation, speech recognition, and time-series prediction. These are just a few examples, and several books have been dedicated to this topic.

Deep Learning

A branch of machine learning called *deep learning* uses artificial neural networks to mimic how the human brain interprets data. These neural networks are made up of layers of connected neurons, or nodes, that can do basic calculations. Deep learning's strength is its capacity to learn from experience and get better without explicit programming. Deep learning is used in various applications, such as image and speech recognition, natural language processing, and autonomous driving.

Large volumes of data are used to train deep learning models. For example, millions of labeled photos may be loaded into a deep learning model for image recognition. The model will begin with simple features like edges and textures at the early levels and progress to more intricate structures like shapes and objects in the later layers. The model progressively extracts details from the photos as it passes the photos deeper into its layers. This hierarchical feature extraction enables the model to recognize and classify images with high accuracy.

Tokenization

Tokenization is the method of dividing text into smaller parts known as tokens. Depending on the tokenization technique employed, these tokens may be words, subwords, or even single characters. Tokenizing the text "Hello, world!" might look like ('Hello' - ',' - 'world' - '!'); this is just one example. This process is essential because it converts unprocessed text into a format that artificial intelligence models can train on.

Tokens are the basic units of input for AI models. Each token is numerically represented, often using embedding techniques that capture the token's semantic meaning in a high-dimensional space. This numerical representation is what the model uses to discover patterns and relationships in the data.

Tokenization helps manage the variability in language. By breaking down text into tokens, models can handle different forms of words, such as plurals or conjugations, more effectively. For example, the word *running* might be tokenized into ("run" - "ning"), allowing the model to recognize the root word "run" across different contexts. The following is a short list of the common tokens used to train AI models.

Word Tokens Each word is treated as a separate token. This is the simplest form of tokenization and can be less effective for languages with complex structure.

Subword Tokens Words are broken down into smaller meaningful units. This method helps handle out-of-vocabulary words better. For example, "unbelievable" might be tokenized into ("un" - "believe" - "able").

Character Tokens Each character is treated as a token. This method is useful for languages with a large number of unique characters or for tasks requiring fine-grained text analysis.

Byte-Pair Encoding (BPE) This is a hybrid approach that merges the most frequently occurring character pairs in a text body. BPE is commonly used in modern NLP models to balance the granularity of tokens.

Tokenization also improves the efficiency of training AI models. Smaller tokens mean that the model can process text in manageable chunks, reducing computational complexity. This is particularly important for large-scale models like transformers, which need to handle vast amounts of data.

Natural Language Processing (NLP)

Natural language processing (NLP) is a component of artificial intelligence that focuses on the interaction between computers and humans through natural language. NLP enables machines to comprehend, interpret, and produce meaningful and practical human language. These tasks include text summarization, sentiment analysis, language translation, and speech recognition. Large volumes of natural language data can be processed and analyzed by AI systems using NLP. This allows AI to automate and enhance a variety of applications such as chatbots, virtual assistants, and language translation services.

Human language is complex and diverse, and this creates problems for NLP. The human language is frequently imprecise, context-dependent, and nuanced, in contrast to structured data. NLP systems combine machine learning algorithms with linguistic rules to overcome these difficulties. Deep learning models assist in using recurrent neural networks (RNNs) to discover patterns and relationships within the data. NLP also uses methods like tokenization, part-of-speech tagging, and named entity identification to aid in the breakdown and classification of text. These developments have greatly increased the precision and effectiveness of NLP applications, improving their dependability and potency.

Another exciting aspect of NLP is its potential to bridge language barriers and facilitate global communication. With the development of sophisticated language models, AI can now translate text and speech between multiple languages with impressive accuracy. This has profound implications for international business, education, and cross-cultural interactions. Additionally, NLP can be used to analyze social media and online content, providing valuable insights into public opinion, trends, and emerging issues.

Foundation Models

Foundation models are large-scale, pretrained models. These models leverage the transformer architecture that allows the model to handle sequential data through mechanisms like self-attention, allowing them to understand and generate complex patterns in text, audio, and video.

There are several different types of foundation models, such as Generative Pretrained Transformer (GPT) and OpenAI's DALL-E, among others. GPT is trained on a large amount

of text-based data, which in turn enables it to generate coherent and contextual information. OpenAI's DALL-E is trained on text-based and image-based data to create AI-generated images. Others like Microsoft Copilot are trained on a mixture of public and private data. These are just a few of the many foundation models available today.

Let's examine the components of a foundation model. GPT stands for Generative Pre-trained Transformer. The generative part is the ability for the model to generate new content from its vast amount of pretrained data. The pretrained component is the data that is fed into the algorithm through machine learning or deep learning. The transformer component of GPT is what allows the foundation model to transform words into serialized data. Transformers can also be used to change serialized data into another form of data, such as images, audio, or even another language. The following are some of the more common transformers used today with AI.

Large Language Models (LLMs)

Large language models (LLMs) are the text-based models that understand written words and that can generate human-like responses. These models have been trained on very diverse datasets from books, articles, websites, literature, and more. This enables the LLM to understand context, nuance, and subtleties in language. They can mimic human communication, and this makes them a valuable tool in applications such as customer service, content creation, quality management, and a vast array of other applications.

They can even be tailored to have personas based on attitude, gender, behavior, and title, just to name a few factors. You could tailor an AI customer service representative to be flexible or rigid or even have the authority of a manager. Another way that LLMs are used is for quality management of calls. Since LLMs have a great grasp on the human language, they can pick up when someone is pleased or agitated, based on language. This sentiment analysis helps call centers with the disposition of calls.

Audio Transformers

The lines blur with audio transformers, because they work in conjunction with LLMs, but they are not specifically tied to text-based models. Audio transformers are used for audio-related tasks, such as speech-to-text and text-to-speech. They can understand spoken words and convert them to sequential data for further processing through LLMs. They can also transform sequential data from the LLMs back to speech with the right inflection, tone, pronunciation, and cadence. One great example of this is Apple's Siri; in recent years it went from a personal assistant to an integral part of Apple's AI ecosystem.

Audio transformers can also perform many other tasks, such as language translation, noise reduction, music generation, and audio classification, and it can even alter human speech. This short list is just a few of the tasks that AI performs for us today.

Video Transformers

Video transformers might seem like a real stretch from audio or text, but they are no different than image transformers. Video transformers either concentrate in an area of the video or they generate an image 30 times a second to create a moving picture. The most impressive and intriguing application of video transformers is deep fakes. They are so impressive and

advanced that the Screen Actors Guild-American Federation of Television and Radio Artists (SAG-AFTRA) went on strike to protest the use of AI in Hollywood.

Although deep fakes are interesting, video transformers are used for a wide array of other applications. Content creation, video editing, surveillance and security, healthcare, sports, and even video-based social media platforms are just a few of the many applications used with video transformers.

Datasets

Data is an essential component of any AI system. The more data that can be pretrained into a model, the better. LLMs and other foundation models need to be pretrained on insane amounts of data. These data sources can be terabytes or even petabytes in size. Remember, pretrained is GPT's middle name.

The old adage for computers, garbage in, garbage out, is true for AI as well. The data source for any AI model needs diversity and quality. The data can be sourced from any of the following channels:

Structured Data This type of data is highly organized and easily searchable. It is typically in the form of relational databases and spreadsheets. Examples include customer information, transaction records, and inventory data.

Unstructured Data Unlike structured data, unstructured data lacks a predefined format. The lack of structure makes it more challenging to process. This category includes text documents, emails, social media posts, images, and videos.

Semi-Structured Data Semi-structured data falls between structured and unstructured data. It includes the elements of both, such as JSON and XML files, which have a loose structure but do not fit neatly into relational databases.

Real-Time Data Real-time data is generated continuously by various sources, such as sensors, IoT devices, and streaming services. This data is crucial for applications that require up-to-date information, such as autonomous vehicles, financial trading systems, and real-time monitoring solutions.

External Data Sources These include publicly available datasets, third-party data providers, and open data platforms. Examples include weather data, economic indicators, and demographic information.

The content of these sources can and should be vast as well. The content can be provided from movie databases, legal databases, scientific journals, research papers, public datasets, government databases, books, music, and the list goes on. Just as a child learns from the time they are born to form structured meaning of their world, AI systems do the same as they are fed diversity and quality of knowledge.

Privacy vs. Security

As AI models are fed with data, they get better with their analysis, prediction, and generation of response. As we use these systems, we contribute to the learning process in a few different

ways. The first is referred to as continuous learning, or online learning, which is when content from the user's input is tokenized and fed into the AI model for pretraining. A great example of this is movie and music streaming services. They present you with the next thing to watch or listen to, based on your likes in comparison to others. Netflix, YouTube, Pandora, and Spotify are just a few content providers that use AI.

If you don't pay for the service, you are the service. That is a great sentiment to live by in today's age of technology. We surrender some, or all of our privacy, when we use services that ingest our participation. Sometimes we voluntarily give up our privacy, and other times we don't even realize we are doing so. Alexa, Apple Home, and Google Home listen to your conversations and spring to attention when you speak to these devices. The AI cloud behind the scenes learns from your input.

There are less obvious ways that information is used and privacy is surrendered, such as your use of your smartwatch or cell phone. These devices can track every breath you take and where you take it. By examining cell phone location data, relationships can be learned. For example, your phone is next to your significant other's between 11 p.m. and 7 a.m. Your phone is in proximity to your colleagues between 8 a.m. and 5 p.m., and from 7 p.m. to 10 p.m. next to your friends. If location services are turned on, patterns of travel can be learned. All this data is potentially available to providers and can be used for AI. There are of course privacy laws that should prevent breaches of your privacy. These are just a few ways we surrender our privacy.

Now that this section made you grab for the tinfoil hat, let's discuss security, which goes hand in hand with privacy. Data security focuses on protecting data from unauthorized access, breaches, and cyberattacks. You might wonder what that has to do with AI? A lot, because AI is pretrained on both material found online and material you enter into it. Threat actors are leveraging AI to uncover and exploit leaked data that may have been inadvertently included in the pretraining datasets of large language models (LLMs).

Human error is often the cause of data leakage where private or proprietary information is inadvertently leaked into LLMs. In a recent incident in 2023, Samsung experienced a data breach when employees inadvertently leaked sensitive information while using ChatGPT. The breach occurred as engineers at Samsung's semiconductor division input confidential data, including source code and internal meeting notes, into ChatGPT to optimize processes and create presentations. This data, once entered into ChatGPT, became part of the AI's training data, making it accessible to OpenAI, the company behind ChatGPT. Samsung has since banned the use of generative AI tools like ChatGPT and is developing its own in-house AI solutions to prevent future leaks. This incident highlights the risks associated with using AI tools in handling sensitive information and underscores the need for stringent data security measures.

Protecting yourself from these security concerns is simple. As an administrator you can create policies that restrict the usage of AI. Microsoft Defender Security supports creating policies that will deter employees from using AI and alert you to the concern. The other tact is to embrace AI, as it is not going away and seems to be entangled in everything. Microsoft offers an AI model called Copilot that keeps your data private by creating a private dataset and segmenting it from public data. In addition, Copilot abides by the security in your

organization. If you ask Copilot a question about the average salary at your organization, it won't dip into Human Resources files, unless you have access to them. This type of AI has been dubbed protected AI or enterprise AI; every vendor has their own buzzword.

Unfortunately, data leakage is not the only inherent problem with AI. Threat actors are using AI tools to help forge emails, research organizations, and even create malware. A threat actor can now formulate an email with these tools and custom-tailor it to your organization. They don't even need to speak your native language! Threat actors can also use tools like ChatGPT to create malware. Typically, mainstream AI tools like ChatGPT, Gemini, and Copilot will censor your requests for illegal activity. However, not all AI tools restrict content, and some sites even promote their use for illegal activity.

Integration

If your organization has not used AI or entertained incorporating it, your organization will not evolve in this new digital age. Competing organizations that do integrate AI will have a competitive advantage. Whether it is creating a sales presentation quicker than your organization, or the allure of AI for your customers, it creates a new competitive race.

So how will your organization stay ahead of the pack? It depends on how your organization chooses to integrate AI. In the early 2010s, business intelligence (BI) became a booming topic for organizations. BI allowed organizations to see relationships in their data, such as top customers, based upon location, time of year, or even the phases of the moon. Today we can use private AI to pretrain our data and extrapolate relationships in data we could not normally see.

Outside of business strategy and AI, Microsoft Copilot allows the summary of Teams meetings, the creation of presentations from Word documents, and meaningful searches of private data with natural language. The marriage of Copilot with the Microsoft 365 suite of products has become a productivity booster.

From an information technology aspect, AI can help you write code, debug code, formulate SQL statements, analyze packet captures, just about anything you can dream up! AI is to network administrators today as the first search engines were to network administrators 25 years ago. We can now use natural language to ask questions and get meaningful answers. One of the coolest tricks with AI is the transformation of information with the use of transformers. You can ask AI to translate pseudocode to PowerShell, Python, or even machine language for an 8080 processor! In Exercise 8.3 you perform this transformation of code.

EXERCISE 8.3

Creating a Script with AI

In the following exercise you will ask ChatGPT to create a PowerShell script from pseudocode. You will then ask it to create a Windows batch script with the same code.

1. Launch your preferred web browser and navigate to ChatGPT by entering `https://chat.openai.com` in the address bar, and press Enter.

EXERCISE 8.3 (continued)

2. Log in to ChatGPT. If you already have an account, click the Log In button. Enter your email address and password, then click Log In. If you don't have an account, click Sign Up and follow the prompts to create one.
3. Compose your request by entering
Write a PowerShell script with the following pseudocode: Loop (print 1, wait, print 0).
4. Review the output and test it in a PowerShell terminal.
5. Now ask ChatGPT to create a Windows batch script from the same code.
6. Review the output and test the code in a batch file.
7. Tell ChatGPT to modify the code, adding counters, then have it transform the code back to PowerShell.

For a bonus, you can have ChatGPT transform the code into other languages, such as C++, Python, or even COBOL!

Limitations

With great power comes great responsibility! This responsibility starts with your use of results from AI. You need to be able to spot the limitations of AI by using some critical thinking applied to your results.

Biases

The answers you get can often exhibit various types of bias. This is typically reflected on the data they are trained on and the design choices made by developers. The choice of algorithm can also be biased due to flawed design or assumptions. The following are some of the biases that you could encounter:

Algorithmic Bias Algorithmic bias occurs when the design of an algorithm favors certain outcomes over other outcomes. This can happen due to the way the algorithm is structured or the assumptions made during its development. For example, if an algorithm is designed to prioritize efficiency over fairness, it might systematically disadvantage certain groups of people. This type of bias can lead to unfair treatment and create inequalities.

Historical Bias Historical bias arises when AI systems are trained on data that reflects past inequalities or prejudices. Since AI models learn from historical data, they can inadvertently inherit and amplify these biases. For instance, if a hiring algorithm is trained on data from a time when certain demographics were underrepresented in certain jobs, it might continue to favor candidates from historically dominant groups, thus reinforcing the status quo.

Data Bias Data bias occurs when the data used to train an AI model is not representative of the real-world population or scenarios it will encounter. This can happen due to sampling errors, data collection methods, or intentional exclusion of certain data points. As a result, the AI system may perform well on the training data but poorly in real-world applications, leading to biased outcomes. For example, facial recognition systems trained predominantly on lighter-skinned individuals may have higher error rates for darker-skinned individuals. This bias made the headlines in 2018, when researchers from MIT and Stanford found this bias in commercial AI systems.

Confirmation Bias Confirmation bias in AI refers to the tendency of algorithms to favor information that confirms preexisting beliefs or hypotheses. This can happen if the training data or the feedback loop reinforces certain patterns while ignoring contradictory evidence. For example, a recommendation system might keep suggesting similar content to users based on their past behavior, thereby limiting their exposure to diverse perspectives and reinforcing their existing preferences.

Emergent Bias Emergent bias arises when an AI system interacts with users or other systems in ways that were not anticipated by its designers. This type of bias can emerge over time as the system learns from new data and adapts to changing environments. For instance, a chatbot might develop biased responses based on the interactions it has with users, especially if it encounters a lot of biased or harmful input. There is more than one story of Internet trolls training a chatbot to say and act in deviant ways. In one instance the chatbot didn't last for more than a few hours.

Evaluation Bias Evaluation bias occurs when the metrics or benchmarks used to assess the performance of an AI system are themselves biased. This can lead to misleading conclusions about the system's effectiveness and fairness. For example, if an AI model is evaluated using a dataset that is not representative of the target population, its performance metrics might not accurately reflect its real-world capabilities, leading to biased decision-making.

Addressing these biases requires a concerted effort to ensure diverse and representative data. Do not trust AI results wholeheartedly and fact-check your results. Knowing that these biases exist and applying some basic critical thinking skills can help you identify bias results.

Hallucinations

AI hallucinations happen when a system produces false or nonsensical information, frequently presenting it with assurance as fact. The intricacy of the task, limits in the training data, or AI's attempt to fill in the blanks with facts that seem true but aren't can contribute to this issue of hallucinations. Consider asking AI to create a pizza recipe. If it has a hallucination, it may recommend an illogical way to adhere cheese on pizza, such as applying glue. This demonstrates how AI can generate strange or inaccurate results when it deviates from its training data or logical thinking, and it is obviously incorrect.

Another instance of hallucinations is when AI provides historical information inaccurately. When AI is questioned about a well-known event, it can make up information.

For example, it might respond that Napoleon Bonaparte won the Battle of Waterloo, which is not true historically. These hallucinations can be problematic, particularly in situations where accuracy is crucial, such as in legal counsel or medical diagnosis. The implications of a hallucination are as follows:

Misinformation AI hallucinations can lead to the dissemination of false information. For instance, if an AI generates incorrect details about a developing emergency, it can mislead the public and hinder effective response efforts.

Healthcare Risks In medical applications, AI hallucinations can be particularly dangerous. An AI might misdiagnose a condition, such as incorrectly identifying a benign lesion as malignant, leading to unnecessary treatments and anxiety for patients.

Legal and Financial In the legal and financial sectors, AI hallucinations can result in incorrect advice or data, potentially causing significant financial losses or legal complications.

Erosion of Trust Frequent AI hallucinations can erode user trust in AI systems. If users repeatedly encounter incorrect or nonsensical outputs, they may become skeptical of the technology's reliability and hesitate to use it in critical applications.

Bias and Discrimination AI hallucinations can also reflect and amplify biases present in the training data. This can lead to discriminatory outcomes, such as biased hiring recommendations or unfair treatment in customer service scenarios.

Operational Inefficiencies In industries like manufacturing or logistics, AI hallucinations can cause operational inefficiencies. For example, an AI might incorrectly predict demand, leading to overproduction or stock shortages.

Educational Impact In educational settings, AI hallucinations can misinform students. If an AI tutor provides incorrect explanations or answers, it can hinder learning and perpetuate misunderstandings.

Security Vulnerabilities AI hallucinations can create security vulnerabilities. For instance, an AI system responsible for monitoring network security might incorrectly identify benign activities as threats or fail to detect actual threats, compromising the security of the system.

Creative Industries In creative fields like writing, art, or music, AI hallucinations can produce outputs that are nonsensical or culturally insensitive, potentially offending audiences or misrepresenting artistic intent.

Customer Service Issues AI-powered customer service bots might provide incorrect information or solutions, frustrating customers and damaging the company's reputation.

Mitigation Strategies

A variety of techniques are used to mitigate AI hallucinations, in turn enhancing the precision and dependability of the AI system. Making sure the training data is of high quality is one of the best strategies. Pretraining the AI system with high-quality data lowers the

possibility of producing erroneous or nonsensical results. It also allows the AI model to learn precise patterns. It is also essential to update the training data on a regular basis to represent correct and most up-to-date information.

Improving the AI model itself is another crucial tactic, but typically it is out of our reach as end users of commercial systems. Other tactics include applying organized data templates, employing sophisticated algorithms that are better able to handle intricate data patterns, and fine-tuning the model with more training on certain tasks or domains. Hallucinations can also be considerably decreased by implementing human-in-the-loop procedures, in which human specialists examine and verify the AI's results. Before any created content is shown to users, this method guarantees that it is correct and trustworthy.

Continuous monitoring and evaluation of AI systems is an essential step for mitigating hallucinations. Implementing validation processes and regularly assessing the AI's performance can help identify and address any issues promptly. By combining these strategies, organizations can significantly reduce the risk of AI hallucinations and ensure that their AI systems are reliable and trustworthy.

Accuracy

Despite AI's immense versatility and power, AI is inherently inaccurate for several reasons. These include the quality and scope of the training data, which may create biases or knowledge gaps as previously discussed. AI models might also have trouble grasping the subtleties and context of the human language, which could result in misunderstandings or mistakes. It is difficult for AI to give accurate responses in every circumstance since real-world scenarios are frequently more complex than the training data. Also, when presented with new or unclear inputs, AI systems may make mistakes because they lack actual knowledge and rely on patterns rather than comprehension. This is evident, for example, if you ever ask AI to create a picture of a cat looking at the moon. You will find that the cat is looking through a telescope like an astronomer, because that is what AI knows when it associates moon and the act of looking or examining.

Considerations

AI has become wildly popular over the past 3 to 5 years. When planning the integration of AI into various applications and processes, be certain to address the considerations of AI adoption by your organization. Appropriate use of AI involves leveraging its capabilities to enhance productivity while respecting privacy and ethics. In this section, we will discuss the considerations for AI in organizations.

Policy

Organizations should establish a comprehensive AI policy that prioritizes ethical considerations, transparency, and accountability. Regular audits and assessments should be conducted to monitor and mitigate potential risks that AI could introduce. Additionally, organizations must plan for continuous learning and improvement in AI, encouraging employees to stay informed about the latest advancements. This should include reviewing the ethical standards of the AI in place.

Some of the typical policies that an organization might adopt are as follows:

Ethical Framework Organizations should develop an ethical framework that guides the development and deployment of AI systems. This framework should include principles such as fairness, accountability, transparency, and respect for user privacy. It should also address issues like bias, discrimination, and the potential for misuse.

Data Governance Establishing robust data governance policies is critical. Organizations must ensure that the data used to train AI models is accurate, relevant, and free from biases. Policies should also cover data security, privacy, and compliance with regulations such as national regulations like General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), just to name a few.

Transparency and Explainability AI systems should be transparent, and their results should be explainable. Organizations need to provide clear information about how AI decisions are made and ensure that users can understand and trust these decisions. This includes documenting the AI models, their training data, and the decision-making processes.

Accountability and Oversight There should be clear lines of accountability for AI systems. Organizations should designate responsible individuals or teams to oversee AI initiatives, conduct regular audits, and ensure compliance with ethical and legal standards. This also involves setting up mechanisms for reporting and addressing any issues or concerns related to AI use.

Continuous Monitoring and Improvement AI systems should be continuously monitored to ensure they perform as intended and do not cause harm. Organizations should implement feedback loops to gather user input and make necessary adjustments. Regular updates and retraining of AI models are essential to maintain their accuracy and relevance.

Employee Training and Awareness Organizations should invest in training programs to educate employees about AI technologies, their potential impacts, and ethical considerations. This helps build a culture of responsible AI use. This training and awareness also empower employees to make informed decisions.

Stakeholder Engagement Engaging with stakeholders, including customers, employees, and regulators, is important for building trust and ensuring that AI policies align with the organization's values and culture. Organizations should seek feedback and collaborate with external experts to refine their AI strategies.

Appropriate Use

Promoting the appropriate use of AI within an organization is important to retain the culture and respect from employees. It is essential to recognize that AI can significantly transform various aspects of business operations. However, it is paramount that the overarching reason to adopt AI is to raise the level of service and not replace people. As an example, a customer service AI-powered chatbot can handle routine inquiries, freeing up human agents to tackle

more complex issues. In finance, AI algorithms can detect fraudulent activities by analyzing transaction patterns in real time, enhancing security and trust.

The appropriate use of AI does not only rest on the shoulders of the organization. The employees of the organization should be responsible for their actions and use of AI for their daily tasks. It's okay to use AI to speed up repetitive tasks, help compile reports, and lend guidance with writing an email. However, the employee engaging in AI needs to treat anything derived from AI as a first draft. The employee should always proof and refine the final product. This collaborative approach ensures that AI complements human capabilities rather than replacing them.

Plagiarism

Plagiarism is a big concern for organizations that publish information, as well as organizations that require end users to provide original works. Academia is one such place that requires students to provide original work to show that they have learned an objective and can synthesize the knowledge. This is similar to instructors requiring you to show mathematics work in math class, rather than just the answer. They want to see that you have learned the concepts and can apply them.

Outside of education, plagiarism is a consideration for corporate organizations. These organizations could suffer legal ramifications if work has been plagiarized by an AI system. The outcome could create an erosion of trust for the organization and carry many other undesirable consequences.

Summary

This chapter focused on the basics of scripting and introduced programming concepts such as scripting types, commenting, branch logic, loops, and use cases for scripting. Scripting is a somewhat newer objective for the CompTIA A+ objectives. It is a worthy skill for an A+ technician.

We also introduced you to the various remote access technologies, highlighting their advantages and weaknesses. In addition, we covered remote assistance tools, which allow a remote user to share their screen with you for support purposes.

Finally we concluded the chapter by discussing artificial intelligence, covering the evolution, components, potential integration, limitations, and considerations for organizations.

Exam Essentials

Understand the basics of scripting. You should understand that a programming language is compiled and a scripting language is interpreted. You should also understand the various elements of a script, such as the comment, interaction of environment variables, branch logic, and loops.

Know the various scripting languages. You should know the advantages and disadvantages of each scripting language and the use cases for scripting, as well as other security and system considerations. In addition, you should know which operating systems each of the various scripting languages operates on. Be familiar with the various script file types that correspond to each scripting language.

Know the various scripting file extensions. The `.bat` and `.cmd` extensions are associated with Windows Batch scripts. The `.ps1` is associated with PowerShell scripts. The `.vbs` extension is associated with VBScript scripts. The `.sh` extension is associated with Bash shell scripts. The `.js` extension is associated with JavaScript scripts. The `.py` extension is associated with Python scripts.

Know the various uses for scripts and their considerations. The various uses for implementing scripts are basic automation, restarting machines, mapping and remapping network drives, installation of applications, automated backups, gathering of information/data, and initiating updates. The considerations are unintentionally introducing malware, inadvertently changing system settings, and browser or system crashes due to mishandling of resources.

Know the various remote access technologies. You should know that RDP is a Microsoft protocol used for remote access. Virtual Network Computing (VNC) is an open standard similar to Microsoft RDP. SSH is an encrypted, secure console protocol. Remote Monitoring and Management (RMM) is a Microsoft protocol for remote administration. Simple Protocol for Independent Computing Environments (SPICE) is an open source protocol used for VDI. In addition, there are third-party screen-sharing and desktop management tools that function similarly and include file-sharing capabilities, as well as videoconferencing software.

Know the various components of artificial intelligence. You should know the various components of AI and understand the training of AI, neural networks, large language models (LLMs), natural language processing, and foundation models.

Understand the limitations and considerations for artificial intelligence. You should understand how biases can flaw the results of AI, as well as what you can do to mitigate the risks of bias. You should also understand how and why AI hallucinations occur and what you can do to prevent propagating wrong information. In addition, you should understand the considerations of policy, appropriate use, and plagiarism with AI.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which statement about scripting languages is true?
 - A. Scripting languages require a compiler.
 - B. Scripting languages are strongly typed.
 - C. Scripting languages are interpreted.
 - D. Scripting languages have good memory management.
2. What level are scripting languages considered?
 - A. High
 - B. Mid
 - C. Intermediate
 - D. Low
3. Which type of variable will allow decimal math?
 - A. Boolean
 - B. Integer
 - C. Floating-point
 - D. String
4. Which environment variable is not inherited?
 - A. System variable
 - B. User variable
 - C. Program variable
 - D. String variable
5. Which statement will load a PowerShell variable `xvar` with a value of 2?
 - A. `xvar = 2`
 - B. `$xvar = 2`
 - C. `xvar = 2;`
 - D. `set /a xvar=2`
6. Which type of loop has a defined beginning and end, and steps from beginning to end?
 - A. `do while` loop
 - B. `while` loop
 - C. `if` statement
 - D. `for` loop

7. Which extension is used with the Windows batch scripting language?
 - A. .vbs
 - B. .js
 - C. .bat
 - D. .py
8. Which scripting language allows for the use of the Component Object Model (COM)?
 - A. PowerShell
 - B. VBScript
 - C. Windows batch script
 - D. JavaScript
9. Which extension is used with the Python scripting language?
 - A. .vbs
 - B. .js
 - C. .bat
 - D. .py
10. Which scripting language is used with Microsoft Azure and Microsoft 365?
 - A. PowerShell
 - B. VBScript
 - C. Windows batch script
 - D. JavaScript
11. Which scripting language is used within web pages to allow for interactive content?
 - A. PowerShell
 - B. Bash
 - C. Windows batch script
 - D. JavaScript
12. Which extension is used with the Bash scripting language?
 - A. .vbs
 - B. .sh
 - C. .bat
 - D. .py
13. What must be done before a Bash script can be executed?
 - A. chown permissions must be set.
 - B. The execute attribute must be set.
 - C. chmod permissions must be set.
 - D. An .sh must be added to the end of the script.

14. Which statement will load a JavaScript variable `mvar` with a value of 8?
- A. `$mvar = 8`
 - B. `mvar = 8`
 - C. `mvar = 8;`
 - D. `set /a mvar=8`
15. Which scripting language has its own preinstalled Integrated Scripting Environment (ISE)?
- A. VBScript
 - B. Bash
 - C. Python
 - D. PowerShell
16. Which line would be used to comment JavaScript code?
- A. `//comment`
 - B. `'comment`
 - C. `REM comment`
 - D. `# comment`
17. Which extension is used with the JavaScript scripting language?
- A. `.js`
 - B. `.sh`
 - C. `.bat`
 - D. `.py`
18. What is the primary goal of artificial intelligence (AI)?
- A. To replace human intelligence
 - B. To simulate or imitate human intelligence
 - C. To create emotional experiences
 - D. To enhance human physical abilities
19. What is the primary function of convolutional neural networks (CNNs)?
- A. Language translation
 - B. Image and video analysis
 - C. Speech recognition
 - D. Game playing
20. What is tokenization in the context of AI?
- A. Dividing text into smaller parts known as tokens
 - B. Converting images into text
 - C. Translating text into another language
 - D. Analyzing audio data

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

You have been assigned to write a PowerShell script that will find other scripts in a user profile directory and all its subdirectories. Which PowerShell variable should you use, since `%UserProfile%` is an environment variable and will not run in PowerShell?



Chapter 9

Safety and Environmental Concerns

**THE FOLLOWING COMPTIA A+ 220-1202
EXAM OBJECTIVES ARE COVERED IN THIS
CHAPTER:**

✓ **4.0 Operational Procedures**

✓ **4.4 Given a scenario, use common safety procedures.**

- Electrostatic discharge (ESD) straps
- ESD mats
- Electrical safety
 - Equipment grounding
- Proper component handling and storage
- Cable management
- Antistatic bags
- Compliance with government regulations
- Personal safety
 - Disconnect power before repairing PC
 - Lifting techniques
 - Fire safety
 - Safety goggles
 - Air filtration mask

✓ **4.5 Summarize environmental impacts and local
environment controls.**

- Material safety data sheet (MSDS)/documentation for handling and disposal
 - Proper battery disposal



- Proper toner disposal
- Proper disposal of other devices and assets
- Temperature, humidity-level awareness, and proper ventilation
 - Location/equipment placement
 - Dust cleanup
 - Compressed air/vacuums
- Power surges, brownouts, and blackouts
 - Uninterruptible power supply (UPS)
 - Surge suppressor

✓ **4.6 Explain the importance of prohibited content/ activity and privacy, licensing, and policy concepts.**

- Incident response
 - Chain of custody
 - Inform management/law enforcement as necessary
 - Copy of drive (data integrity and preservation)
 - Incident documentation
 - Order of volatility
- Licensing/digital rights management (DRM)/end-user license agreement (EULA)
 - Valid licenses
 - Perpetual license agreement
 - Personal-use license vs. corporate-use license
 - Open-source license
- Non-disclosure agreement (NDA)/mutual non-disclosure agreement (MNDA)
- Regulated data
 - Credit card payment information
 - Personal government-issued information



- PII
- Healthcare data
- Data retention requirements
- Acceptable use policy (AUP)
- Regulatory and business compliance requirements
 - Splash screens



In this chapter, we start by talking about safety, which includes your safety and the safety of your coworkers, as well as environmental concerns. Observing proper safety procedures can help prevent injury to you or to others.

Our discussion about the environment is two-sided. The environment affects computers (via things like dust, sunlight, and water), but computers can also potentially harm the environment. We'll consider both sides as we move through this chapter.

Next, we will cover some legal aspects of operational procedures. These include licensing of software, protection of personally identifiable information, and incident response.



Exam Objective 4.4, "Given a scenario, use common safety procedures," requires you to know the various safety procedures. You should know how electrostatic discharge (ESD) can be controlled, proper component handling and storage, compliance with government regulations, and personal safety.

Understanding Safety Procedures

The proliferation of computers in today's society has created numerous jobs for technicians. Presumably that's why you're reading this book: you want to get your CompTIA A+ certification. Many others, who don't fix computers professionally, like to tinker with them as a hobby. Years ago, only the most expert users dared to crack the case on a computer. Often, repairing the system meant using a soldering iron. Today, thanks to the cheap parts, computer repair is not quite as involved. Regardless of your skill or intent, if you're going to be inside a computer, you always need to be aware of safety issues. There's no sense in getting yourself hurt or killed.

As a provider of a hands-on service (repairing, maintaining, or upgrading someone's computer), you need to be aware of some general safety tips, because if you are not careful, you could harm yourself or the equipment. Clients expect you to solve their problems, not make them worse by injuring yourself or those around you. In this section, we'll talk about identifying safety hazards and creating a safe working environment.

Identifying Potential Safety Hazards

Anything can be a potential safety hazard, right? Okay, maybe that statement is a bit too paranoid, but there *are* many things, both human-made and environmental, that can cause safety problems when you're working with and around computers.

Perhaps the most important aspect of computers you should be aware of is that not only do they *use* electricity, but they also *store* electrical charge after they're turned off. This makes the power supply and older *cathode-ray tube (CRT)* monitor pretty much off-limits to anyone but a repairperson trained specifically for those devices. In addition, the computer's processor and various parts of the printer run at extremely high temperatures, and you can get burned if you try to handle them immediately after they've been in operation.

Those are just two general safety measures that should concern you. There are plenty more. When discussing safety issues with regard to PCs, let's break them down into four general areas:

- Computer components
- Electrostatic discharge
- Electromagnetic interference
- Natural elements

Computer Components

As mentioned earlier, computers use electricity. And as you're probably aware, electricity can hurt or kill you. The first rule when working inside a computer is always to make sure that it's powered off. If you have to open the computer to inspect or replace parts (as you will with most repairs), be sure to turn off the machine before you begin. Leaving it plugged in is usually fine, and many times it is actually preferred because it grounds the equipment and can help prevent electrostatic discharge.

There's one exception to the power-off rule: You don't have to power off the computer when working with hot-swappable parts, which are designed to be unplugged and plugged back in when the computer is on. Most of these components have an externally accessible interface (such as USB devices or hot-swappable hard drives), so you don't need to crack the computer case.

The Power Supply

Do not take the issue of safety and electricity lightly. Removing the power supply from its external casing can be dangerous. The current flowing through the power supply normally follows a complete circuit; when your body breaks the circuit, your body becomes part of that circuit. Getting inside the power supply is the most dangerous thing you can do as an untrained technician.

The two biggest dangers with power supplies are burning or electrocuting yourself. These risks usually go hand in hand. If you touch a bare wire that is carrying current, you could get electrocuted. A large-enough current passing through you can cause severe burns. It can also cause your heart to stop, your muscles to seize, and your brain to stop functioning. In short, it can kill you. Electricity always finds the best path to ground. And because the human body is basically a bag of saltwater (an excellent conductor of electricity), electricity will use us as a conductor if we are grounded.

Although it is possible to open a power supply to work on it, doing so is *not* recommended. Power supplies contain several capacitors that can hold *lethal* charges *long after they have been unplugged!* It is extremely dangerous to open the case of a power supply.

Besides, power supplies are relatively inexpensive and are considered *field replaceable units (FRUs)*. It would probably cost less to replace one than to try to fix it—and much safer.

In the late 1990s, a few mass computer manufacturers experimented with using open power supplies in their computers to save money. We don't know if any deaths occurred because of such incompetence, but it was definitely a very bad idea.

Current vs. Voltage: Which Is More Dangerous?

When talking about power and safety, you will almost always hear the saying, "It's not the volts that kill you; it's the amps." That's mostly true. However, an explanation is in order.

The number of volts in a power source represents its potential to do work. But volts don't do anything by themselves. Current (amperage, or amps) is the force behind the work done by electricity. Here's an analogy to help explain this concept: Say you have two boulders. One weighs 10 pounds and another weighs 100 pounds, and each is 100 feet off the ground. If you drop them, which one will do more work? The obvious answer is the 100-pound boulder. They both have the same potential to do work (100 feet of travel), but the 100-pound boulder has more mass and thus more force. Voltage is analogous to the distance the boulder is from the ground, and amperage is analogous to the mass of the boulder.

This is why you can produce static electricity on the order of 50,000 volts and not electrocute yourself. Even though this electricity has a great potential for work, it does very little work because the amperage is so low. This also explains why you can weld metal with a 110- or 220-volt welder. Welders use only 17 to 45 volts at the tip of the welder, but they also use anywhere from 55 to 590 amps!

If you ever have to work on a power supply, for safety's sake you should discharge all capacitors within it. To do this, connect a resistor across the leads of the capacitor with a rating of 3 watts or more and a resistance of 100 ohms (Ω) per volt. For example, to discharge a 225-volt capacitor, you would use a 22.5k Ω resistor ($225V \times 100\Omega = 22,500\Omega$, or 22.5k Ω). You can also purchase a tool, known as a *grounding pen* or *discharge pen*, that will discharge any size capacitor. The tools look similar to a pen and have a wire that is connected to ground.

The Monitor

Other than the power supply, the most dangerous component to try to repair is an older style *cathode-ray tube (CRT)* computer monitor—specifically. In fact, we recommend that you *not* try to repair monitors of any kind.

The CRT monitor is a lift hazard due to the weight of these monitors. A CRT monitor can retain a high-voltage charge for hours after it's been turned off. The CRT monitor is also under a constant vacuum, which makes them implode if the glass is broken. This could cause sharp glass to scatter in all directions.

Even though we recommend not repairing monitors, the A+ exam may test your knowledge of the safety practices to use if you ever need to do so. If you have to open a CRT monitor, you must first discharge the high-voltage charge on it by using a *high-voltage probe*. This probe has a very large needle, a gauge that indicates volts, and a wire with an alligator clip. Attach the alligator clip to a ground (usually the round pin on the power cord). Slip the probe needle underneath the high-voltage cup on the monitor. You will see the gauge spike to around 15,000 volts and slowly reduce to 0 (zero). When it reaches 0, you may remove the high-voltage probe and service the high-voltage components of the monitor.



Do not use an ESD strap when discharging the CRT monitor. Doing so can lead to a fatal electric shock.

Working with *liquid crystal display (LCD)* monitors—or any device with a fluorescent or LCD backlight—presents a unique safety challenge. These types of devices require an *inverter*, which provides the high-voltage, high-frequency energy needed to power the backlight.

The inverter is a small circuit board installed behind the LCD panel that takes DC power and converts (inverts) it for the backlight. If you've ever seen a laptop or handheld device with a flickering screen or perpetual dimness, it was likely an inverter problem. Inverters store energy even when their power source is cut off, so they have the potential to discharge that energy if you mess with them. Be careful! Newer LED displays use LED light to back-light the display, but there is no clear way to distinguish between LCD backlit and LED backlit monitors, so use caution.

The Case

One component that people frequently overlook is the case. Cases are generally made of metal, and some computer cases have very sharp edges inside, so be careful when handling them. You can cut yourself by jamming your fingers between the case and the frame when you try to force the case back on. Also of particular interest are drive bays. Countless technicians have scraped or cut their hands on drive bays when trying in vain to plug a drive cable into the motherboard. You can cover particularly sharp edges with duct tape. Just make sure that you're covering only metal and nothing with electrical components on it.

The Printer

If you've ever attempted to repair a printer, you might have thought that a little monster was inside, hiding all the screws from you. Besides missing screws, here are some things to watch out for when repairing printers:

- When handling a toner cartridge from a laser printer or page printer, do not turn it upside down. You will find yourself spending more time cleaning the printer and the surrounding area than fixing the printer.
- Do not put any objects into the feeding system (in an attempt to clear the path) when the printer is running.

- Laser printers generate a laser that is hazardous to your eyes. Do not look directly into the source of the laser.
- If it's an inkjet printer, do not try to blow in the ink cartridge to clear a clogged opening—that is, unless you like the taste of ink.
- Some parts of a laser printer (such as the *EP cartridge*, which is an *image drum*) will be damaged if you touch them. Your skin produces oils and has a small surface layer of dead skin cells. These substances can collect on the delicate surface of the image drum (EP cartridge) and cause malfunctions. Bottom line: keep your fingers out of places where they don't belong.
- Laser printers use very high-voltage power sources to charge internal components, which can cause severe injuries.
- Laser printers can get extremely hot in order to fuse the toner to the page. Don't burn yourself on internal components. The fuser can get extremely hot!



Using an egg carton (or other container with small compartments) is a great way to store and keep track of screws that you take out of a device when you're working on it. If it's an impromptu repair and you don't have a container with compartments, you can use a piece of paper with tape that is looped. This allows you to also write down what the screws were removed from.

When working with printers, we follow some pretty simple guidelines. If there's a messed-up setting, paper jam, or ink or toner problem, we will fix it. If it's something other than that, we call a certified printer repairperson. The inner workings of printers can get pretty complex, and it's best to call someone trained to make those types of repairs.

The Keyboard and Mouse

Okay, we know that you're thinking, "What danger could a keyboard or mouse pose?" We admit that not much danger is associated with these components, but there are a couple of safety concerns that you should always keep in mind:

- If your mouse has a cord, it can catch on something, causing items to fall off a desk, or it might even be long enough that someone can trip over it. So, make sure that it's safely out of the way.
- You could short-circuit your keyboard if you accidentally spill liquid into it. Keyboards generally don't function well with half a can of cola in their innards.

Electrostatic Discharge

So far, we've talked about how electricity can hurt people, but it can also pose safety issues for computer components. One of the biggest concerns for components is *electrostatic discharge (ESD)*. For the most part, ESD won't do serious damage to a person other than

provide a little shock. But little amounts of ESD can cause serious damage to computer components, and that damage can manifest itself by causing computers to hang, reboot, or fail to boot at all. ESD happens when two objects of dissimilar charge come into contact with each other. The two objects exchange electrons in order to standardize the electrostatic charge between them. This charge can, and often does, damage electronic components.



CPU chips and memory chips are particularly sensitive to ESD. Be extremely cautious when handling them.

When you shuffle your feet across the floor and shock your best friend on the ear, you are discharging static electricity into their ear. The lowest static voltage transfer that you can feel is around 3,000 volts; it doesn't electrocute you because there is extremely little current. A static transfer that you can *see* is at least 10,000 volts! Just by sitting in a chair, you can generate around 100 volts of static electricity. Walking around wearing synthetic materials can generate around 1,000 volts. You can easily generate around 20,000 volts simply by dragging your smooth-soled shoes across a carpeted floor in the winter. (Actually, it doesn't have to be winter. This voltage can occur in any room with very low humidity—like a heated room in wintertime.)



Relative humidity has a significant impact on the electricity you generate. Walking around can generate 1,500 volts at 65–90 percent relative humidity, but it can produce 35,000 volts if the relative humidity is in the 10–25 percent range.

It makes sense that these thousands of volts can damage computer components. However, a component can be damaged with less than 300 volts! This means that if a small charge is built up in your body, you could damage a component without realizing it.




Do you have long hair or (gasp!) have to wear a tie when fixing computers? Tie it back. Letting long hair or dangling cloth get inside an open computer case is asking for trouble, because both are notorious for carrying and conducting static electricity.

The good news is that there are measures you can implement to help contain the effects of ESD. The first and easiest item to implement is the antistatic wrist strap, also referred to as an ESD strap. We will look at the antistatic wrist strap, as well as other ESD prevention tools, in the following sections.

Antistatic Wrist Straps

To use an ESD *antistatic strap*, you attach one end to an earth ground (typically, the ground pin on an extension cord) or the computer case and wrap the other end around your wrist. This strap grounds your body and keeps it at a zero charge. Figure 9.1 shows the proper way to attach an antistatic strap. Several varieties of wrist straps are available. The strap shown in Figure 9.1 uses an alligator clip and is attached to the computer case itself, whereas others use a banana clip that attaches to a grounding coupler.

FIGURE 9.1 ESD strap using an alligator clip

Real World Scenario

ESD Symptoms

Symptoms of ESD damage may be subtle, but they can be detected. One of the authors relates this experience:

“When I think of ESD, I always think of the same instance. Several years ago, I was working on an Apple Macintosh. This computer seemed to have a mind of its own. I would troubleshoot it, find the defective component, and replace it. The problem was that as soon as I replaced the component, it failed. I thought maybe the power supply was frying the boards, so I replaced both at the same time, but to no avail.

“I was about to send the computer off to Apple when I realized that it was winter. Normally this would not be a factor, but winters where I live are extremely dry. Dry air promotes static electricity. At first I thought my problem couldn’t be that simple, but I was at the end of my rope. So, when I received my next set of new parts, I grounded myself with an anti-static strap for the time it took to install the components, and prayed while I turned on the power. Success! The components worked as they should, and a new advocate of ESD prevention was born.”

In order for an antistatic wrist strap to work properly, the computer must be plugged in but turned off. When the computer is plugged in, it is grounded through the power cord. When you attach yourself to it with the wrist strap, you are grounded through the power cord as well. If the computer is not plugged in, there is no ground, and any excess electricity on you will just discharge into the case, which is not good.



An ESD strap is a device that is specially designed to bleed electrical charges away *safely*. It uses a 1 megohm resistor to bleed the charge away slowly. A simple wire wrapped around your wrist will not work correctly, and you could be electrocuted!




Never wear an ESD strap if you're working inside a monitor or inside a power supply. If you wear one while working on the inside of these components, you increase the chance of getting a lethal shock.

ESD Antistatic Mats

It is possible to damage a device by simply laying it on a benchtop. Therefore, you should have an ESD *antistatic mat* in addition to an ESD strap. An ESD mat drains excess charge away from any item coming in contact with it (see Figure 9.2). ESD mats are also sold as mouse/keyboard pads to prevent ESD charges from interfering with the operation of the computer. Many wrist straps can be connected to the mat, thus causing the technician and any equipment in contact with the mat to be at the same electrical potential and eliminating ESD. There are even ESD bootstraps and ESD floor mats, which are used to keep the technician's entire body at the same potential.

FIGURE 9.2 Proper use of an ESD antistatic mat



Antistatic Bags for Parts

Antistatic bags, shown in Figure 9.3, are important tools to have at your disposal when servicing electronic components because they protect the sensitive electronic devices from stray static charges. By design, the static charges collect on the outside of these silver or pink bags, rather than on the electronic components.

FIGURE 9.3 An antistatic component bag



Unlike antistatic mats, antistatic bags do not “drain” the charges away, and they should never be used in place of an antistatic mat. These bags conduct a uniform charge, because the surface is somewhat conductive.

You can obtain the bags from several sources. The most direct way to acquire antistatic bags is to go to an electronics supply store and purchase them in bulk. Most supply stores have several sizes available. Perhaps the easiest way to obtain them, however, is simply to hold onto the ones that come your way. That is, when you purchase any new component, it usually comes in an antistatic bag. After you have installed the component, keep the bag. It may take you a while to gather a collection of bags if you take this approach, but eventually you will have a fairly large assortment.

Self-Grounding

We recommend that you include a grounding strap in your toolkit so that you’re never without it. But we also realize that things happen and you might find yourself in a situation where you don’t have your strap or an ESD mat. In such cases, you should *self-ground*.

Self-grounding is not as effective as using proper anti-ESD gear, but it makes up for that with its simplicity. To self-ground, make sure the computer is turned off but plugged in. Then touch an exposed (but not hot or sharp) metal part of the case. This will drain electrical charge from you. Better yet is if you can maintain constant contact with that metal part.

That should keep you at the same bias as the case. Yes, it can be rather challenging to work inside a computer one-handed, but it can be done.

Additional Methods


Another preventive measure that you can take is to maintain the relative humidity at around 50 percent. Don't increase the humidity too far—to the point where moisture begins to condense on the equipment. It is best to check with the manufacturer of the equipment that you are protecting to find the optimal humidity. Also, use antistatic spray, which is available commercially, to reduce static buildup on clothing and carpets.



If you don't have any antistatic spray, you can always use the "Downy solution." In a spray bottle, combine one part water with one part liquid fabric softener. Mist areas such as carpet and clothing that cause problems. If used regularly, it will keep static away—and keep your office smelling nice.

Vendors have methods of protecting components in transit from manufacture to installation. They press the pins of integrated circuits (ICs) into antistatic foam to keep all the pins at the same potential, as shown in Figure 9.4. In addition, most circuit boards are shipped in antistatic bags, as discussed earlier.

FIGURE 9.4 Antistatic foam



Antistatic foam looks a lot like Styrofoam. However, there are huge differences between the two. While antistatic foam helps reduce the transfer of electricity, Styrofoam does not. Styrofoam holds a charge on its surface quite easily. Have you ever tried to get some of those small packing "peanuts" off your hands? Be careful to not mix the two up, lest you fry your components.

At the very least, you should be mindful of the dangers of ESD and take steps to reduce its effects. Beyond that, you should educate yourself about those effects so that you know when ESD is becoming a major problem.

Electromagnetic Interference

When compared to the other dangers that we've discussed in this chapter, *electromagnetic interference (EMI)*, also known as *radio-frequency interference (RFI)* when it's in the same frequency range as radio waves, is by far the least dangerous. EMI really poses no threats to you in terms of bodily harm. What it can do is make your equipment or network malfunction.

EMI is an unwanted disturbance caused by electromagnetic radiation generated by another source. In other words, some of your electrical equipment may interfere with other equipment. Here are some common sources of interference:

Network Devices The popularity of wireless networking devices has introduced the possibility of interference. Some of the most popular wireless networking standards, 802.11b/g/n/ax/be, use the 2.4 GHz range for transmissions. Bluetooth devices happen to use the same frequency. In theory, they won't interfere with each other because they use different modulation techniques. In practice, however, interference between the two types of devices can happen.

Magnets Magnets work by generating an electromagnetic field. It might make sense, then, that this field could cause electromagnetic interference. For the most part, you don't need to worry about this unless you have huge magnets at work. Do note, however, that many motors use magnets, which can cause interference. For example, one of our friends used to have his computer on the opposite side of a wall from his refrigerator. Whenever the compressor kicked in, his monitor display would become wavy and unreadable. It was time to move his home office. Another common culprit is a desk fan. Put a desk fan next to a monitor and turn on the fan. What happens to the display? It will become wavy. This is another example of EMI.

Cordless Phones Cordless phones can operate at a variety of frequencies. Some of the more common frequencies are 900 MHz, 1.9 GHz, 2.4 GHz, and 5.8 GHz. Many of these are common ranges for computer equipment to operate in as well.

Microwave Ovens Microwave ovens are convenient devices for heating food and beverages. The radiation they generate is typically in the 2.45 GHz range, although it can vary slightly. If a microwave is being used near your computer, you'll often see a distorted display, just as if a fan or motor were being run next to your computer. You may also experience interference with wireless network communications.



Copper wires are susceptible to EMI. Fiber-optic cables, which use light to transmit data, are not susceptible to EMI.

Natural Elements

Computers should always be operated in cool environments and away from direct sunlight and water sources. This is also true when you're working on computers. We know that heat is an enemy of electrical components. Dirt and dust act as great insulators, trapping heat inside components. When components run hotter than they should, they have a greater chance of breaking down faster.

It pretty much should go without saying, but we'll say it anyway: water and electricity don't mix. Keep liquids away from computers. If you need your morning coffee while fixing a PC, make sure that the coffee cup has a tight and secure lid.



Real World Scenario

Water and Servers Don't Mix

This situation happened at a company where one of the authors used to work. The building needed some roof repairs. Repairs went on for several days, and then the weekend came. It just so happened that the area on which they were working was over the server room. That weekend was a particularly rainy one, and, of course, no one was in the office over the weekend.

Monday morning came, and the IT staff arrived to find that the server room was partially flooded. Rain had come in through weaknesses in the roof caused by the maintenance and had flooded through the drop ceiling and into the server room. Nearly half a million dollars' worth of equipment was ruined.

Although this is an extreme example, it illustrates an important point: always be aware of the environment in which you're working, and be alert to potential sources of problems for your computer equipment.

Creating a Safe Workplace

Benjamin Franklin was quoted as saying, "An ounce of prevention is worth a pound of cure." That sage advice applies to a lot in life and certainly to computer safety. Knowing how to work with and handle computer equipment properly is a good start. It's also important to institutionalize and spread the knowledge, and to make sure that your company has the proper policies and procedures in place to ensure everyone's safety.

Moving Computer Equipment

We've already talked about some of the hazards posed by computer parts. Many times it's the more mundane tasks that get us, though, such as moving stuff around. One of the most

common ways that IT employees get hurt is by moving equipment in an improper way. Changing the location of computers is a task often completed by IT personnel. You can avoid injury by moving things the right way.

To ensure your personal safety, here are some important techniques to consider before moving equipment:

- The first thing always to check for is that it's unplugged. There's nothing worse (and potentially more dangerous) than getting yanked because you're still tethered.
- Securely tie the power cord to the device or remove it altogether, if possible.
- Remove any loose jewelry and secure long hair or neckties.
- Lift with your legs, not your back. Bend at the knees when picking something up, not at the waist.
- Do not twist when lifting.
- Maintain the natural curves of the back and spine when lifting.
- Keep objects close to your body and at waist level.
- Push rather than pull, if possible.

The muscles in the lower back aren't nearly as strong as those in the legs or other parts of the body. Whenever lifting, you want to reduce the strain on those lower-back muscles as much as possible. If you want, use a back belt or brace to help you maintain the proper position while lifting.



CRT monitors can be heavy. (Thank goodness for flat screens!) When lifting and carrying any type of monitor, always keep the glass face toward your body. The front of the monitor is the heaviest part, and you want the heavy part closest to your body in order to reduce strain on your muscles.

If you believe that the load is too much for you to carry, don't try to pick it up. Get assistance from a coworker. Another great idea is to use a cart. It will save you trips if you have multiple items to move, and it saves you the stress of carrying components.

If you do use a cart to move the equipment, make sure that you do not overload the cart. Know the cart's weight limitation and estimate the weight of the equipment you will be hauling. Most small, commercial service carts will hold around 100–200 pounds. If you're moving a battery backup unit that requires two people to lift, you may be pushing the limitations of the cart. Also make sure the load is not top heavy. Always place the heaviest items on the lower shelves of a cart.

When moving loads, always be aware of your surrounding environment. Before you move, scout out the path to see whether there are any *trip hazards* or other safety concerns, such as spills, stairs, uneven floors (or ripped carpet), tight turns, or narrow doorways.

Using Appropriate Repair Tools

A big part of creating a safe working environment is having the right tools available for the job. There's no sense implementing a sledgehammer solution to a ball-peen hammer

problem. Using the wrong tool might not help fix the problem, and it could very possibly hurt you or the computer in the process.

Most of the time, computers can be opened and devices removed with nothing more than a simple screwdriver. But if you do a lot of work on PCs, you'll definitely want to have additional tools on hand.

Computer toolkits are readily available on the Internet or at any electronics store. They come in versions from inexpensive (under \$10) kits that have around 10 pieces to kits that cost several hundred dollars and have more tools than you will probably ever need. Figure 9.5 shows an example of a basic 13-piece *PC toolkit*. All of these tools come in a handy zippered case, so it's hard to lose them.

FIGURE 9.5 PC toolkit



Figure 9.5 shows the following tools, from left to right:

- Sockets, ranging from 1/4" and 1/2"
- Screwdriver bits: #1 and #1 Phillips tips, #1 and #2 flat tip, T10, T15, T20 Torx tips, and a 1/4" drive adapter
- Hex keys
- ESD strap
- Screw container

- Screwdriver
- 1/8" flat screwdriver
- Tweezers
- Long-nose pliers
- A three-claw retriever
- A #2 Phillips extension
- An integrated circuit (IC) extractor
- Wire strippers

A favorite of ours is the three-claw retriever, because screws like to fall and hide in tiny places. While most of these tools are incredibly useful, an IC extractor probably won't be. In today's environment, it's rare to find an IC that you can extract, much less find a reason to extract one.

The following sections look at some of the tools of the PC troubleshooting trade.

Screwdrivers

Every PC technician worth their weight in pocket protectors needs to have a *screwdriver*—at least one. There are three major categories of screwdrivers: flat-blade, Phillips, and Torx. In addition, there are devices that look like screwdrivers, except that they have a hex-shaped indented head on them. They're called *hex drivers*, and they belong to the screwdriver family.

When picking a screwdriver, always keep in mind that you want to match the size of the screwdriver head to the size of the screw. Using a screwdriver that's too small will cause it to spin inside the head of the screw, stripping the head of the screw and making it useless. If the screwdriver is too large, on the other hand, you won't be able to get the head in far enough to generate any torque to loosen the screw. Of course, if the screwdriver is way too big, it won't even fit inside the screw head at all. Common sizes for Phillips-head screws are 000, 00, 0, 1, 2, and 3. When you are dealing with Torx screws, the two most common sizes are T-10 and T-15.



When tightening screws, you don't need to make them so tight that they could survive the vibrations of an atmospheric reentry. Snug is fine. Making them too tight can cause problems when loosening them, which could cause you (or someone else not so strong) to strip the head. Using an electric screwdriver is fine, if you have one. The only problem with them is that they tend to be larger than manual screwdrivers and can be difficult to get inside a case.



Using magnetic-tipped screwdrivers is not recommended. Many computer disks contain magnetically coded information, and the magnetic tip of a screwdriver could cause a problem. Keep a retrieving tool handy instead, just in case you drop a screw.

Antistatic Wrist Straps

We've already talked about these, but they are important, so we'll mention them again. An antistatic wrist strap is essential to any PC technician's arsenal. They don't typically come with smaller PC toolkits, but you should always have one or two handy.

Other Useful Tools

PC techs also commonly carry the following tools:

Pliers Pliers are useful for a variety of tasks, especially gripping something. Long-nose or needle-nose pliers extend your reach.

Wire Cutters Wire cutters come in a variety of forms but are primarily used for cutting cables. It's not likely that you'll need any sort of heavy-duty metal cutters.

Wire Strippers If you are making your own network cables or fixing them, having a cable stripper (and crimper) is essential.

Mirrors Mirrors are handy inside tight spaces. Many techs like to use a dentist-style mirror because of its compact size and good reach.

Flashlight Never underestimate the utility of a good flashlight. You never know what your lighting situation will be like when you're at a repair site. Smaller flashlights with good output are great to have because they can fit into tight spaces.

Compressed Air For as much as computers and dust don't get along, it sure seems like they are attracted to each other. In all seriousness, computer components are powered by electricity, which causes the components to have a slight electrical charge. Dust is also electrically charged, so it's attracted to computer components. Compressed air can help you clean off components, especially in hard-to-reach places.



Be judicious about your use of compressed air. Often, you will find yourself just blowing the dust from one part of a computer to another.

Multimeter If you're having power issues, a multimeter is an invaluable tool. It measures electrical current, voltage, and resistance. You'll also hear of voltmeters, and while the two have somewhat different functions, both of them can be used to troubleshoot power problems. You can use a voltmeter to determine whether a computer's power supply is producing the right amount of current for the devices that depend on it.

Air Filter Masks We've mentioned dust, and you will occasionally encounter large amounts of it when working inside computers. If you're using compressed air to blow dust out, it could easily get into your eyes or lungs. You don't want to breathe these airborne particles, and in some environments (like machine shops), you don't know what

types of hazardous materials are mixed in with the airborne particles. You don't need a big army-grade gas mask, but a small respirator will help considerably.

Safety Goggles Your eyesight should be a primary concern when working with any type of tools, because accidents happen. You should always wear protective glasses that cover all sides of your eyes. When you think of *safety goggles*, you probably think of the science lab rubberized goggles from high school chemistry class. Safety goggles can be as comfortable as a pair of glasses, if you purchase approved safety glasses.

Fire Safety

Repairing a computer isn't often the cause of an electrical fire. However, you should know how to extinguish such a fire properly. Four major classes of fire extinguishers are available, one for each type of flammable substance: A, for wood and paper fires; B, for flammable liquids; C, for electrical fires; and D (metal powder or NaCl [salt]), for flammable metals, such as phosphorus and sodium.


The most popular type of fire extinguisher today is the multipurpose, or ABC-rated, extinguisher. It contains a dry chemical powder (e.g., sodium bicarbonate, monoammonium phosphate) that smothers the fire and cools it at the same time. For electrical fires (which may be related to a shorted-out wire in a power supply), make sure the fire extinguisher will work for Class C fires. If you don't have an extinguisher that is specifically rated for electrical fires (Class C), you can use an ABC-rated extinguisher.

The topic of electrical fire safety is a very broad subject. The best prevention method for electrical fires is to follow building codes. If building codes are not followed, your organization could be fined. Every state and locality has a different building code that you should reference with new structures and alterations to existing structures. In addition to state and local building codes, you also need to reference national building codes related to fire prevention. One organization that supplies reference to these codes is the National Fire Protection Association (NFPA) at www.nfpa.org/for-professionals/codes-and-standards. The U.S. Fire Administration (www.usfa.fema.gov/prevention) is a government organization that provides general training and prevention material.


Creating a Safe Work Environment

We've already talked about some work environment issues. For example, don't put a computer next to the breakroom sink, and keep computers out of direct sunlight (even if the desk location is great). A few other things to watch out for are trip hazards, atmospheric conditions, and high-voltage areas.

Cables are a common cause of tripping. If at all possible, run cables through drop ceilings or through conduits to keep them out of the way. If you need to lay a cable through a trafficked area, use a floor cable guard to keep the cables in place and safe from crushing. Floor guards come in a variety of lengths and sizes (for just a few cables or for a lot of cables). Figure 9.6 shows a cable guard.

FIGURE 9.6 Floor cable guard

Another useful tool to keep cables under control is a *cable tie* (see Figure 9.7). It's simply a plastic tie that holds two or more cables together. Cable ties come in different sizes and colors, so you're bound to find one that suits your needs.

FIGURE 9.7 Cable ties



In a pinch, and without a floor cable guard, you can use tape, such as gaffing tape, to secure your cables to the floor. This is recommended only as a temporary fix for two reasons. First, it's not much less of a trip hazard than just having the cables run across the floor. Second, gaffing tape doesn't protect the cables from being crushed if people step on them or heavy objects are moved over them.

Exercise 9.1 is a simple procedure that you can modify and use as needed. Its purpose is to illustrate common office hazards that you may not have realized were there.

EXERCISE 9.1

Finding Trip Hazards

1. Walk around the server room and count how many cables are lying on the floor.
2. Walk around the client areas and see how many cables are lying on the floor or are exposed underneath cubicles.

Maybe you're fortunate and don't find any, but odds are that you found at least one area with exposed cables that should not be exposed. You can reapply this exercise for other dangerous items, such as exposed wires and exposed sharp edges.

Atmospheric conditions you need to be aware of include areas with high static electricity or excessive humidity. Being aware of these conditions is especially important for preventing electrostatic discharge, as we've already discussed.

Finally, be aware of high-voltage areas. Computers do need electricity to run, but only in measured amounts. Running or fixing computers in high-voltage areas can cause problems for the electrical components and problems for you if something should go wrong.

Implementing Safety Policies and Procedures

The *Occupational Safety and Health Act* states that every working American has the right to a safe and healthy work environment. To enforce the act, the *Occupational Safety and Health Administration* (OSHA) was formed. OSHA covers all private-sector employees and U.S. Postal Service workers. Public-sector employees are covered by state programs, and federal employees are covered under a presidential executive order. In a nutshell, OSHA requires employers to "provide a workplace that is free of recognized dangers and hazards."

There are three overarching criteria to a safe work environment:

- The company and its employees have identified all significant hazards in the work setting.
- Preventive measures have been taken to address each significant hazard.
- The company and its employees understand how to respond to accidents or near-miss accidents if or when they occur.

Next, let's explore specific responsibilities and how to create a safe work environment plan.



Always ensure that your company's safety policies and procedures comply with all government regulations.

Employer and Employee Responsibilities

Maintaining workplace safety is the responsibility of employers as well as employees. Here are some of the important responsibilities of employers:

- Provide properly maintained tools and equipment.
- Provide a warning system, such as codes or labels, to warn employees of potential hazards or dangerous chemicals.
- Post the OSHA poster in a prominent location.
- Keep records of workplace injuries or illnesses.
- Continuously examine workplace conditions to ensure OSHA compliance.

It's also the responsibility of the employee to help maintain a safe work environment. Specifically, employees are charged with the following tasks:

- Read and understand OSHA posters.
- Follow all employer-implemented health and safety rules and safe work practices.
- Use all required protective gear and equipment.
- Report hazardous conditions to the employer.
- Report hazardous conditions that the employer does not correct to OSHA.

As you can see, employers and employees need to work together to keep the workplace safe. It is illegal for an employee to be punished in any way for exercising their rights under the Occupational Safety and Health Act.



Real World Scenario

Play It Safe with Common Sense

When you're repairing a PC, do not leave it unattended. Someone could walk into the room and inadvertently bump the machine, causing a failure. Worse, they could step on pieces that may be lying around and get hurt. It is also not a good idea to work on the PC alone. If you're injured, someone should be around to help. Finally, if you're fatigued, you may find it difficult to concentrate and focus on what you are doing. The most important safety measure to remember when you are repairing PCs is to pay close attention to what you are doing.

Safety Plans

We recommend that your company create and follow a workplace safety plan. Having a safety plan can help avoid accidents that result in lost productivity, equipment damage, and employee injury or death.

A good safety plan should include the following elements:

- A written document that states, among other things, who is responsible for implementing and managing the plan
- Systematic periodic inspections to identify workplace hazards
- Procedures for eliminating hazards once identified
- Processes for investigating the cause of accidents, injuries, or illnesses
- A safety and health training program specific to the job duties performed
- A system for employees to communicate safety or health concerns without fear of reprisal
- A system to ensure that employees comply with safety and health rules
- A system to maintain safety and health records, including steps taken to implement accident-prevention initiatives

It might seem like a laundry list of items to consider, but a good safety program needs to be holistic in nature to be effective.

Many companies also incorporate rules against drug or alcohol use in their safety and health plans. Specifically, employees are not allowed to come to work if under the influence of alcohol or illegal drugs. Employees who do come to work under the influence may be subject to disciplinary action up to and including termination of employment.

After your safety plan has been created, you need to ensure that all employees receive necessary training. Have each employee sign a form at the end of the training to signify that they attended, and keep the forms in a central location (such as with or near the official safety policy). In addition to the training record, you should make available and keep records of the following:

- Safety improvement suggestion form
- Accident and near-accident reporting form
- Injury and illness log
- Safety inspection checklist
- Hazard removal form
- Material safety data sheets

Safety rules and regulations will work only if they have the broad support of management from the top down. Everyone in the organization needs to buy into the plan; otherwise, it won't be a success. Make sure that everyone understands the importance of a safe work environment, and make sure that the culture of the company supports safety in the workplace.

Incident Management

Accidents happen. Hopefully, they don't happen too often, but we know that they do. Details on how to handle accidents are a key part of any safety plan so that when an accident does happen, you and your coworkers know what to do. Good plans should include steps for handling a situation as well as reporting an incident. We will cover *incident response* in more detail later in this chapter. Two major classifications of accidents are environmental and human.

When related to computers, environmental accidents typically come in one of two forms: electricity or water. Too much electricity is bad for computer components. If lightning is striking in your area, you run a major risk of frying computer parts. Even if you have a surge protector, you could still be at risk.

The best bet in a lightning storm is to power off your equipment and unplug it from outlets. Make the lightning have to come inside a window and hit your computer directly in order to fry it.



Those cheap \$10 surge suppressors will fry right along with your computer. And don't be fooled—most power strips do *not* protect against power surges.

Water is obviously also bad for computer components. If there is water in the area and you believe that it will come in contact with your computers, it's best to get the machines powered off as quickly as possible. If components are not powered on but get wet, they may still work after thoroughly drying out. But if they're on when they get wet, they're likely cooked. Water + electronic components = bad. Water + electronic components + electricity = *really* bad.

Many server rooms have raised floors. Although this serves several purposes, one is that equipment stored on the raised floor is less susceptible to water damage if flooding occurs.

Human nature dictates that we are not infallible, so, of course, we're going to make mistakes and have accidents. The key is to minimize the damage caused when an accident happens.

If a chemical spill occurs, make sure that the area gets cordoned off as soon as possible. Then clean up the spill. The specific procedure on how to do that depends on the chemical, and that information can be found on material safety data sheets (MSDSs). Depending on the severity of the spill or the chemical released, you may also need to contact the local authorities. Again, the MSDS should have related information. We cover MSDSs in more detail later in the chapter.

Physical accidents are more worrisome. People can trip on wires and fall, cut or burn themselves repairing computers, and incur a variety of other injuries. Computer components can be replaced, but that's not always true of human parts (and it's certainly not true of lives). The first thing to keep in mind is always to be careful and use common sense. If you're trying to work inside a computer case and you see sharp metal edges inside the case, see whether the metal (or component on which you are working) can be moved to another location until you finish. Before you stick your hand into an area, make sure that nothing is hot or could cut you.

When an accident does happen (or almost happens), be sure to report it. Many companies pay for workers' compensation insurance. If you're injured on the job, you're required to

report the incident, and you might also get temporary payments if you are unable to work because of the accident. Also, if the accident was anything but minor, seek medical attention. Just as victims in auto accidents might not feel pain for a day or two, victims in other physical accidents might be in the same position. If you never report the accident, insurance companies may find it less plausible that your suffering was work related.



Exam Objective 4.5, “Summarize environmental impacts and local environment controls,” requires you to know the impact of the physical environment on electronic equipment. This includes power issues and how you can protect your equipment from common power issues. In addition, you should know how to properly dispose of electronics.

Understanding Environmental Controls

It is estimated that more than 25 percent of all the lead (a poisonous substance) in landfills today comes from consumer electronics components. Because consumer electronics (televisions, cell phones, laptops, stereos) contain hazardous substances, many states require that they be disposed of as *hazardous waste*. Computers are no exception. CRT monitors contain several carcinogens and phosphors as well as mercury and lead. The computer itself may contain several lubricants and chemicals as well as lead. Printers contain plastics and chemicals, such as those in toners and inks, which are also hazardous. All of these items should be disposed of properly.

Remember all those 386 and 486 computers that came out in the late 1980s and are now considered antiques? Maybe you don’t, but there were millions of them. Where did they all go? Is there an Old Computers Home somewhere that is using these computer systems for good purposes, or are they lying in a junkyard somewhere? Or could it be that some folks just cannot let go and have a stash of old computer systems and computer parts in the dark depths of their basements? Regardless of where they are today, all of those old components have one thing in common: They are hazardous to the environment.

On the flip side, the environment is also hazardous to our computers. We’ve already talked about how water and computers don’t mix well, and that’s just the beginning. Temperature, humidity, and air quality can have dramatic effects on a computer’s performance. And we know that computers require electricity; too much or too little can be a problem.

With all these potential issues, you might find yourself wondering, “Can’t we all just get along?” In this section, we will talk about how to make our computers and the environment coexist as peacefully as possible.

Managing the Physical Environment

Some of our computers sit in the same dark, dusty corner for their entire lives. Other computers are carried around, thrown into bags, and occasionally dropped. Either way, the

physical environment in which our computers exist can have a major effect on how long they last. It's smart to inspect the physical environment periodically in order to ensure that there are no working hazards. Routinely cleaning components will also extend their useful life, and so will ensuring that the power supplying them is maintained.

Maintaining Power

As electronics, computers need a power source. Laptops can free you from your power cord leash for a while, but only temporarily. Power is something that we often take for granted until we lose it, and then we twiddle our thumbs and wonder what people did before the Internet. Most people realize that having too much power (a *power surge*) is a bad thing because it can fry electronic components. Having too little power, such as when a *blackout* occurs, can also wreak havoc on electrical circuits.

Obviously, if we lose power, the equipment stops working. We have all experienced blackouts, but there are many other electrical problems we can encounter that will affect our network equipment and interrupt operations:

Sags An electrical sag is created when the line voltage dips below the nominal operating voltage. Normally the voltage should stay consistent at 90 percent of the voltage specification. An electrical sag is considered an event if the voltage drops within 1/2 second to 3 seconds of a power cycle.

Brownout An electrical sag (undervoltage event) where the line voltage dips below the nominal operating voltage for more than 3 seconds of a power cycle. Brownouts are commonly caused by large upstream draws on electrical demand.

Spike Power spikes are created when the line voltage increases over the nominal operating voltage. A voltage spike can be tens of thousands of volts, but the event only lasts 1 to 2 nanoseconds.

Surges Power surges are also created when the line voltage spikes over the nominal operating voltage. However, the difference between a spike and a surge is that a surge can last over 3 nanoseconds.

Power strips come in all shapes and sizes and are convenient for plugging multiple devices into one wall outlet. Most of them even have an on/off switch so that you can turn all the devices on or off at the same time. Figure 9.8 shows a simple power strip.

Don't make the mistake of thinking that power strips will protect you from electrical surges, though. If you get a strong power surge through one of these \$10 devices, the strip and everything plugged into it can be fried. Some people like to call power strips "surge protectors" or "surge suppressors," but power strips do nothing to protect against or suppress surges.

Devices that actually attempt to keep power surges at bay are called *surge protectors*. They often look similar to a power strip, so it's easy to mistake them for each other, but protectors are more expensive, usually starting in the \$25 range. Surge protectors have a fuse inside them that is designed to blow if it receives too much current and not to transfer the current to the devices

plugged into it. Surge protectors may also have plug-ins for RJ-11 (phone), RJ-45 (Ethernet), Universal Serial Bus (USB), and F-connector (coaxial cable) connectors.

FIGURE 9.8 A simple power strip



Figure 9.9 shows a surge protector, which doesn't look too different from a simple power strip. The key is to read the packaging and the labels on the product. Make sure that the device will protect your electronics from electrical surges. There is usually a printed specification of 200V to 500V. This indicates how much of a surge in voltage the surge protector can handle.

FIGURE 9.9 A surge protector



The best device for power protection is called an *uninterruptible power supply (UPS)*. These devices can be as small as a brick, like the one shown in Figure 9.10, or as large as an entire server rack. Some just have a few indicator lights, while others have LCD displays that show status and menus and that come with their own management software.

FIGURE 9.10 An uninterruptible power supply



Inside the UPS are one or more batteries and fuses. Much like a surge suppressor, a UPS is designed to protect everything that's plugged into it from power surges. UPSs are also designed to protect against *power sags* and even power outages. Energy is stored in the batteries, and if the power fails, the batteries can power the computer for a period of time so that the administrator can then safely power it down. Many UPSs and operating systems will also work together to power down a system automatically (and safely) or switch it to UPS power. These types of devices may be overkill for Uncle Bob's machine at home, but they're critically important fixtures in server rooms.

UPSs can accommodate several different devices; the number depends on the size and power rating. The model shown in Figure 9.11 has four plugs for battery backup and surge protection, and another four outlets for surge protection only. Two of each of the four outlets are controlled by a master switch on the unit.

The UPS should be checked periodically as part of the preventive maintenance routine to make sure its battery is operational. Most UPSs have a test button you can press to simulate a power outage. You will find that batteries wear out over time, and you should replace the battery in the UPS every couple of years to keep the UPS dependable.

FIGURE 9.11 The back of a UPS

Power blackouts are generally easy to detect. Power sags without a complete loss, called a *brownout*, are also very damaging to electrical components but often go unnoticed. Uninterruptable power supplies can prevent this type of damage. A UPS provides enough time to safely shut down equipment during a power outage. It protects electronic devices and offers immediate emergency power in case of failure.

Managing the Environment

Sometimes we can't help how clean—or unclean—our environments are. A computer in an auto body shop is going to face dangers that one in a receptionist's office won't. Still, there are things that you can do to help keep your systems clean and working well. We're going to break these concepts down into two parts. First, we'll look at common issues you should be aware of, and then we'll discuss proper cleaning methods.

Avoiding Common Problems

In a nutshell, water and other liquids, dirt, dust, unreliable power sources, and heat and humidity aren't good for electronic components. Inspect your environment to eliminate as many of these risks as possible. Leaving your laptop running outside in a rainstorm? Not such a good idea. (Been there, done that.)

Computers in manufacturing plants are particularly susceptible to environmental hazards. One technician reported a situation with a computer that had been used on the manufacturing floor of a large equipment manufacturer. The computer and keyboard were covered with a black substance that would not come off. (It was later revealed to be a combination

of paint mist and molybdenum grease.) There was so much diesel fume residue in the power supply fan that it would barely turn. The insides and components were covered with a thin, greasy layer of muck. To top it all off, the computer smelled terrible!

Despite all this, the computer still functioned. However, it was prone to reboot itself every now and again. The solution was (as you may have guessed) to clean every component thoroughly and replace the power supply. The muck on the components was able to conduct a small current. Sometimes, that current would go where it wasn't wanted and zap—a reboot. In addition, the power supply fan is supposed to partially cool the inside of the computer. In this computer, the fan was detrimental to the computer because it got its cooling air from the shop floor, which contained diesel fumes, paint fumes, and other chemical fumes. Needless to say, those fumes aren't good for computer components.

Computers and humans have similar tolerances to heat and cold, although computers like the cold better than we do. In general, anything comfortable to us is comfortable to a computer. They don't, however, require food or drink (except maybe a few RAM chips now and again)—keep those away from the computer.



It's bad practice to eat, drink, or smoke around your computer. Smoke particles contain tar that can get inside the computer and cause problems similar to those just described.

Computers need lots of clean moving air to keep them functioning. One way to ensure that the environment has the least possible effect on your computer is always to leave the blanks in the empty slots on the back of your box. These pieces of metal are designed to keep dirt, dust, and other foreign matter out of the inside of the computer. They also maintain proper airflow within the case to ensure that the computer does not overheat.

You can also purchase computer enclosures to keep the dust out—just make sure that they allow for proper air ventilation. Many times these devices use air filters in much the same way a furnace or a car engine does.

Cleaning Systems

The cleanliness of a computer is extremely important. Buildup of dust, dirt, and oils can prevent the various mechanical parts of a computer from operating. Cleaning them with the right cleaning compounds is equally important. Using the wrong compounds can leave residue behind that is more harmful than the dirt that you are trying to remove.

Most computer cases and monitor cases can be cleaned by using mild soapy water on a clean, lint-free cloth. Do *not* use any kind of solvent-based cleaner on monitor screens, because doing so can cause discoloration and damage to the screen surface. Most often, a simple dusting with a damp cloth (moistened with water) will suffice. Make sure that the power is off before you put anything wet near a computer. Dampen (don't soak) a cloth in mild soap solution and wipe the dirt and dust from the case. Then wipe the moisture from the case with a dry, lint-free cloth. Anything with a plastic or metal case can be cleaned in this manner.



Don't drip liquid into any vent holes on equipment. Monitors in particular often have vent holes in their tops.

Additionally, if you spill anything on a keyboard, you can clean it by soaking it in distilled, *demineralized water* and then drying it off. The extra minerals and impurities have been removed from this type of water, so it will not leave any traces of residue that might interfere with the proper operation of the keyboard after cleaning. The same holds true for the keyboard's cable and its connector.

The electronic connectors of computer equipment, on the other hand, should never touch water. Instead, use a swab moistened in distilled, *denatured isopropyl alcohol* (also known as electronics or contact cleaner and found in electronics stores) to clean contacts. Doing so will take oxidation off the copper contacts.

Finally, the best way to remove dust and debris from the inside of the computer is to use compressed air (not a vacuum). Compressed air can be more easily directed and doesn't easily produce ESD damage, as a vacuum could. Simply blow the dust from inside the computer by using a stream of compressed air. Make sure to do this outside so that you don't blow dust all over your work area or yourself. Also be sure to wear safety goggles and use an air mask. If you need to use a vacuum, a nonstatic *computer vacuum* that is specially made for cleaning computer components is recommended. Their nozzles are grounded to prevent ESD from damaging the components of the computer.

One unique challenge when cleaning printers is spilled toner. It sticks to everything and should not be inhaled—it's a carcinogen. Use an electronics vacuum that is designed specifically to pick up toner. A typical vacuum's filter isn't fine enough to catch all the particles, so the toner may be circulated into the air. Normal electronics vacuums may melt the toner instead of picking it up.



If you get toner on your clothes, use a magnet to get it out. (Toner is half iron.)

Table 9.1 summarizes the most common cleaning tools and their uses.

TABLE 9.1 Computer cleaning tools

Tool	Purpose
Computer vacuum	Sucking up dust and small particles.
Mild soap and water	Cleaning external computer and monitor cases.
Demineralized water	Cleaning keyboards or other devices that have contact points that are not metal.
Denatured isopropyl alcohol	Cleaning metal contacts, such as those on expansion cards.

Tool	Purpose
Monitor wipes	Cleaning monitor screens. Do <i>not</i> use window cleaner.
Lint-free cloth	Wiping down anything. Don't use a cloth that will leave lint or other residue behind.
Compressed air	Blowing dust or other particles out of hard-to-reach areas.

Periodically cleaning equipment is one of the easiest ways to prevent costly repairs, but it's also one of the most overlooked tasks. We're often too busy solving urgent crises to deal with these types of tasks. If possible, block out some time every week for the sole purpose of cleaning your equipment.

Handling and Disposing of Computer Equipment

Each piece of computer equipment that you purchase offers a manual, usually found online. Detailed instructions on the proper handling and use of that component can be found in the manual. In addition, many manuals give information on how to open the device for maintenance or on whether you should even open the device at all.

If you have the luxury of having paper manuals, don't throw them away. Keep a drawer of a file cabinet specifically for hardware manuals (and keep it organized). You can always look up information on the Internet as well, but having paper manuals on hand is useful for two reasons:

- You may need to fix something when Internet access isn't readily available. (Router problems, anyone?)
- Some companies are required to keep hardware documentation, in the event of an audit (such as for ISO 9000-compliant organizations).

In the following sections, we'll cover two topics: using safety documentation and following safety and disposal procedures.



There are several different ISO specifications that an organization can be certified for by the International Organization for Standardization (ISO). Each different ISO certification corresponds to a specific set of policies and procedures that an organization must adhere to. When a company is certified, it means they meet or exceed following the standardized policies and procedures for their organization.

Using Safety Documentation

In addition to your product manuals, another place to find safety information is in material safety data sheets (MSDSs). MSDSs include information such as physical product data (boiling point, melting point, flash point, and so forth), potential health risks, storage and disposal recommendations, and spill/leak procedures. With this information, technicians and emergency personnel know how to handle the product as well as respond in the event of an emergency.

MSDSs are typically associated with hazardous chemicals. Indeed, chemicals do not ship without them. MSDSs are not intended for consumer use; rather, they're made for employees or emergency workers who are consistently exposed to the risks of the particular product.

The U.S. *Occupational Safety and Health Administration (OSHA)* mandates MSDSs only for the following products:

- Products that meet OSHA's definition of *hazardous* (they pose a physical or health hazard)
- Products that are "known to be present in the workplace in such a manner that employees may be exposed under normal conditions of use or in a foreseeable emergency"

One of the interesting things about MSDSs is that OSHA does not require companies to distribute them to consumers. Most companies will be happy to distribute one for their products, but they are under no obligation to do so.

If employees are working with materials that have MSDSs, those employees are required by OSHA to have "ready access" to MSDSs. This means that employees need to be able to get to the sheets without having to fetch a key, contact a supervisor, or submit a procedure request. Remember the file cabinet drawer that you have for the hardware manuals? MSDSs should also be kept readily accessible. Exercise 9.2 helps you find your MSDSs and get familiar enough with them to find critical information.

EXERCISE 9.2

Finding MSDSs

1. Locate the MSDSs in your workplace. You might have to ask a manager. (Do you even have them?)
2. Find an MSDS for a product you're interested in.
3. Does the MSDS list any potential health effects for this item? What are they?
4. What is the proper disposal procedure for this item?

It's not likely that you're going to memorize or need to memorize everything on an MSDS. The key things are to know where to find the MSDS and how to find information in it quickly. If you have a spill of a potentially dangerous chemical, the last thing you need to do is spend your time figuring out how to handle the spill without causing injury to yourself or others.

At this point, you might stop to think for a second, “Do computers really come with hazardous chemicals? Do I really need an MSDS?” Consider this as an example: oxygen. Hardly a dangerous chemical, considering we need to breathe it to live, right? In the atmosphere, oxygen is at 21 percent concentration. At 100 percent concentration, oxygen is highly flammable and can even spontaneously ignite some organic materials. In that sense, and in the eyes of OSHA, nearly everything can be a dangerous chemical.



If you are interested in searching for free MSDSs, several websites are available, such as www.msds.com. Many manufacturers of components also provide MSDSs on their websites.

The sections within an MSDS are the same regardless of the product, but the information inside each section changes. Here is a truncated sample MSDS for ammonium hydrogen sulfate:

**** MATERIAL SAFETY DATA SHEET ****

Ammonium Hydrogen Sulfate

90009

**** SECTION 1—CHEMICAL PRODUCT AND COMPANY IDENTIFICATION ****

MSDS Name: Ammonium Hydrogen Sulfate

Catalog Numbers:

A/5400

Synonyms:

Sulfuric acid, monoammonium salt; Acid ammonium sulfate; Ammonium acid sulfate.

**** SECTION 2—COMPOSITION, INFORMATION ON INGREDIENTS ****

CAS# Chemical Name % EINECS#

7803-63-6 Ammonium hydrogen sulfate 100 % 232-265-5

Hazard Symbols: C

Risk Phrases: 34

**** SECTION 3—HAZARDS IDENTIFICATION ****

EMERGENCY OVERVIEW

Causes burns. Corrosive. Hygroscopic (absorbs moisture from the air).

Potential Health Effects

Skin:

Causes skin burns.

Ingestion:

May cause severe gastrointestinal tract irritation with nausea, vomiting, and possible burns.

Inhalation:

Causes severe irritation of upper respiratory tract with coughing, burns, breathing difficulty, and possible coma.

**** SECTION 4—FIRST-AID MEASURES ****

Skin:

Get medical aid immediately. Immediately flush skin with plenty of water for at least 15 minutes while removing contaminated clothing and shoes.

Ingestion:

Do not induce vomiting. If victim is conscious and alert, give 2–4 cupfuls of milk or water. Never give anything by mouth to an unconscious person. Get medical aid immediately.

Inhalation:

Get medical aid immediately. Remove from exposure and move to fresh air immediately.

If not breathing, give artificial respiration. If breathing is difficult, give oxygen.

**** SECTION 5—FIREFIGHTING MEASURES ****

**** SECTION 6—ACCIDENTAL RELEASE MEASURES ****

General Information: Use proper personal protective equipment as indicated in Section 8.

**** SECTION 7—HANDLING and STORAGE ****

Handling:

Wash thoroughly after handling. Wash hands before eating. Use only in a well-ventilated area. Do not get in eyes, on skin, or on clothing. Do not ingest or inhale.

Storage:

Store in a cool, dry place. Keep container closed when not in use.

**** SECTION 8—EXPOSURE CONTROLS, PERSONAL PROTECTION ****

Engineering Controls:

Use adequate general or local exhaust ventilation to keep airborne concentrations below the permissible exposure limits.

Respirators:

Follow the OSHA respirator regulations found in 29 CFR 1910.134 or European Standard EN 149. Always use a NIOSH or European Standard EN 149 approved respirator when necessary.

**** SECTION 9—PHYSICAL AND CHEMICAL PROPERTIES ****

Physical State: Solid

Color: White

Odor: Not available

**** SECTION 10—STABILITY AND REACTIVITY ****

Chemical Stability:

Stable under normal temperatures and pressures.

Conditions to Avoid:

Incompatible materials, dust generation, exposure to moist air or water.

**** SECTION 11—TOXICOLOGICAL INFORMATION ****

RTECS#:

CAS# 7803-63-6: BS4400500

**** SECTION 12—ECOLOGICAL INFORMATION ****

**** SECTION 13—DISPOSAL CONSIDERATIONS ****

Products which are considered hazardous for supply are classified as Special Waste, and the disposal of such chemicals is covered by regulations which may vary according to location. Contact a specialist disposal company or the local waste regulator for advice. Empty containers must be decontaminated before returning for recycling.

**** SECTION 14—TRANSPORT INFORMATION ****

**** SECTION 15—REGULATORY INFORMATION ****

European/International Regulations

European Labeling in Accordance with EC Directives

Hazard Symbols: C

Risk Phrases:

R 34 Causes burns.

Safety Phrases:

S 26 In case of contact with eyes, rinse immediately with plenty of water and seek medical advice. S 28 After contact with skin, wash immediately with...

**** SECTION 16—ADDITIONAL INFORMATION ****

MSDS Creation Date: 6/23/2004 Revision #0 Date: Original.

Following Proper Disposal Procedures

It is relatively easy to put old components away, thinking that you might be able to put them to good use again someday, but doing so is not realistic. Most computers are obsolete as soon as you buy them. And if you have not used them recently, your old computer components will more than likely never be used again.

We recycle cans, plastic, and newspaper, so why not recycle computer equipment? The problem is that most computers contain small amounts of hazardous substances. Some countries are exploring the option of recycling electrical machines, but not all have enacted appropriate measures to enforce their proper disposal.



Some countries are ahead of others on the recycling issue. For example, in 2013, the United Kingdom introduced the Waste Electrical and Electronic Equipment (WEEE) Regulations, which obligates manufacturers to provide disposal methods for items they manufacture. Consumers can insist that manufacturers take back dangerous waste and electrical and electronic equipment. Many local governments also provide e-waste programs and disposal centers.

Regardless of manufacturer or community programs, we can take proactive steps, as consumers and caretakers of our environment, to promote the proper disposal of computer equipment:

- Check with the manufacturer. Some manufacturers will take back outdated equipment for parts (and may even pay you for them).
- Properly dispose of solvents or cleaners (as well as their containers) used with computers at a local hazardous waste disposal facility.
- Disassemble the machine and reuse the parts that are good.
- Check out businesses that can melt down the components for the lead or gold plating.
- Contact the Environmental Protection Agency (EPA) for a list of local or regional waste disposal sites that accept used computer equipment. The EPA's web address is www.epa.gov.
- Check with the EPA or at www.msds.com to see if what you are disposing has an MSDS. These sheets contain information about the toxicity of a product and whether it can simply be disposed of as trash. They also contain lethal-dose information.
- Check with local nonprofit or education organizations that may be interested in using the equipment.

Check out the Internet for possible waste disposal sites. Table 9.2 lists a few websites that we came across that deal with the disposal of used computer equipment. A quick web search will likely locate some in your area.

TABLE 9.2 Computer recycling websites

Site name	Web address
The Home Depot	www.homedepot.com/c/electronics_recycling_programs
Staples	https://www.staples.com/stores/recycling
U.S. EPA	www.epa.gov/recycle/electronics-donation-and-recycling
Repowered	https://getrepowered.org



The National Conference of State Legislatures details each state's legislation regarding electronics waste. The site also details state programs to collect electronics waste. For more information, visit:

www.ncsl.org/environment-and-natural-resources/electronic-waste-recycling

Following the general rule of thumb of recycling your computer components and consumables is a good way to go. In the following sections, we'll look at four classifications of computer-related components and the proper disposal procedures for each.

Batteries

The EPA estimates that more than 350 million batteries are purchased annually in the United States. One can only imagine what the worldwide figure is. Batteries contain several heavy metals and other toxic ingredients, including alkaline, mercury, lead acid, nickel cadmium, and nickel metal hydride.



Never burn a battery to destroy it. Doing so will cause the battery to explode and could result in serious injury.

When batteries are thrown away and deposited into landfills, the heavy metals inside them will find their way into the ground. From there, they can pollute water sources and eventually find their way into the supply of drinking water. In 1996, the United States passed the *Mercury-Containing and Rechargeable Battery Management Act* (aka the *Battery Act*) with two goals: to phase out the use of mercury in disposable batteries and to provide collection methods and recycling procedures for batteries.



Several countries around the world have battery-recycling programs. Information on battery recycling in the United States can be found at www.call2recycle.org. You can find information on the Battery Act at www.epa.gov/rcra/mercury-containing-and-rechargeable-battery-management-act-public-law-104-142

Five types of batteries are most commonly associated with computers and handheld electronic devices: alkaline, nickel-cadmium (NiCd), nickel-metal hydride (NiMH), lithium-ion (Li-ion), and *button cell*.

Alkaline Batteries *Alkaline batteries* have been incredibly popular portable batteries for several decades. Before 1984, one of the major ingredients in this type of battery was mercury, which is highly toxic to the environment. In 1984, battery companies began reducing the mercury levels in batteries, and in 1996, mercury was outlawed in alkaline batteries in the United States. Still, it's strongly recommended that you recycle these batteries at a recycling center. Although newer alkaline batteries contain less mercury than their predecessors, they are still made of metals and other toxins that contaminate the air and soil.

Nickel-Cadmium *Nickel-cadmium (NiCd)* is a popular format for rechargeable batteries. As their name indicates, they contain high levels of nickel and cadmium. Although nickel is only semi-toxic, cadmium is highly toxic. These types of batteries are categorized by the EPA as hazardous waste and should be recycled.

Nickel-Metal Hydride (NiMH) and Lithium-Ion (Li-ion) Laptop batteries are commonly made with *Nickel-Metal Hydride (NiMH)* and *lithium ion (Li-ion)*. Unlike the previous types of batteries that we have discussed, these are not considered hazardous waste, and there are no regulations on recycling them. However, these batteries do contain elements that can be recycled, so it's still a good idea to go that route.

Button Cell These batteries are so named because they look like a button. They're commonly used in calculators and watches as well as portable computers. They often contain mercury and silver (and are environmental hazards due to the mercury) and need to be recycled.

You may have noticed a theme regarding the disposal of batteries: recycling. Many people just throw batteries in the trash and don't think twice about it. However, there are several laws in the United States that require the recycling of many types of batteries, and recycling does indeed help keep the environment clean. For a list of recycling centers in your area search the Internet.



If you're ever exposed to the electrolyte (the inside "juice") of the battery, immediately flush the exposed area with water. If it gets on your eye, immediately contact a physician and wash the eye for 15 minutes.

Display Devices

Computer monitors (CRT monitors, not LCDs) are big and bulky, so what do you do when it's time to get rid of them? As previously mentioned, monitors contain capacitors that are capable of retaining a lethal electric charge after the monitors have been unplugged. You wouldn't want anyone to set off the charge accidentally and die. But what we didn't mention earlier, which is important now, is that most CRT monitors contain high amounts of lead. Most monitors contain several pounds of lead, in fact. Lead is very dangerous to humans and the environment and must be dealt with carefully. Other harmful elements found in CRTs include arsenic, beryllium, cadmium, chromium, mercury, nickel, and zinc.

If you have to dispose of a monitor, contact a computer-recycling firm. It's best to let professional recyclers handle the monitor for you.



Real World Scenario

How Not to Dispose of Your Monitors

This story comes from the technical support division of a now-defunct major computer manufacturer, which used a lot of computers at its own facility. At one time, the company had as many as 500 technicians working the phones. So you can imagine that they burned out a lot of equipment.

Here's how they disposed of dead monitors. An IT staff member would take the monitor out to the dumpster and bring along a sledgehammer. Setting the monitor on its back, he

would take one good swing at the glass panel with the hammer to shatter the screen. (This was done, by policy, to ensure that no one would want to go out to the dumpster and try to salvage the dead monitor.) After spraying glass everywhere, he picked up the monitor and threw it in the dumpster.

One employee observed that it probably wasn't good to be spreading glass all over the parking lot by shattering monitors. That advice was taken, and the sledgehammer was retired. Instead, an IT staff member would use a permanent black marker and draw all over the screen (again, so no one would want to try to salvage it), and again, it was thrown in the dumpster.

In our enlightened state today (as opposed to the mid-1990s), we can see how this was not a good plan for disposing of broken monitors. In fact, many states today have laws prohibiting the disposal of computer monitors in trash bins. This is a good law because, with the amount of harmful elements in monitors, they're every bit the environmental hazard that batteries are.

Laser Printer Toner Disposal

Toner cartridges should be recycled as well. PC recycling centers will take old toner cartridges and properly dispose of them. The toner itself is a carcinogen, and the cartridges can contain heavy metals that are bad for the environment.

Toner cartridges are valuable to companies that refurbish and refill these cartridges. It's actually big business to refill these expensive cartridges, and it's environmentally responsible. Most toner is the same for all types and models of laser printers, so a toner refurbishing center will refill and test these cartridges. Then these companies sell the cartridges for a fraction of the price of new toner cartridges. If a new toner cartridge is installed, the old toner cartridge is boxed up and sent back. Even if your organization doesn't contract with one of these services, they would be happy to take the old cartridges off your hands and keep them out of the trash.

Mobile Devices

Cell phones and tablets are considered disposable units, with the average life expectancy of two to four years. Their popularity has outpaced mobile computing and the desktop computer market. These mobile devices are extremely small and fit neatly into the trash. However, mobile devices contain the same toxic metals and chemical compositions as their larger cousins. Every mobile device has a circuit board that contains lead, batteries that contain other heavy metals, and bodies that contain plastics.

These devices should be recycled responsibly. Many big-box retailers have an anonymous recycling drop bin, where you can recycle a mobile device. There are even automated kiosks where you can get money for a defective device. You simply tell it what is wrong with the device, the device's make and model, and its condition. The kiosk will then offer you a few

dollars to recycle it. The company that owns and operates the kiosk refurbishes the device and resells it for a fraction of the cost of a new one.



Before recycling a mobile device, make sure the device is wiped of your personal data with the factory reset function. You never know where the device will end up. It could be refurbished and resold locally or sold for parts to a nefarious market.

Chemical Solvents and Cans

Nearly every chemical solvent that you encounter will have a corresponding MSDS. On the MSDS for a chemical, you will find a section detailing the proper methods for disposing of it. Chemical solvents were not designed to be released into the environment, because they could cause significant harm to living organisms if they're ingested. If in doubt, contact a local hazardous materials handler to find out the best way to dispose of a particular chemical solvent.

Cans are generally made from aluminum or other metals, which are not biodegradable. It's best always to recycle these materials. If the cans were used to hold a chemical solvent or otherwise hazardous material, contact a hazardous materials disposal center instead of a recycling center.



Always be sure that you are following all applicable laws and regulations when disposing of computer equipment.



Exam Objective 4.6, "Explain the importance of prohibited content/activity and privacy, licensing, and policy concepts," requires you to know how to perform incident response, the various licensing/digital rights management (DRM), end-user license agreement (EULA), non-disclosure agreement (NDA)/mutual nondisclosure agreement (MNDA), acceptable use policy (AUP), and regulatory and business compliance requirements.

Understanding Policies, Licensing, and Privacy

Many of the operational procedures that we've discussed up to this point have been about safety—yours, your computer equipment's, and the environment's. We've also touched on regulations, as in always be sure to comply with local government regulations. In this section, we focus more on the legal side of things. Not understanding legal requirements is not

a justifiable defense in a court of law. Considering that IT professionals often deal with software licensing and personally identifiable information, or sometimes encounter prohibited activity or have to deal with a security incident, you should understand the general principles related to these concepts.

Dealing with Prohibited Content/Activity

This is a situation that no one really wants to deal with, but it happens more often than we would care to admit: a computer you are fixing has content on it that is inappropriate or illegal, or you see someone on your network performing an action that is against policy or laws. How you respond in such a situation can have a significant bearing on your career, the other people involved, and, depending on the situation, the well-being of your organization. The key to dealing with *prohibited content* or activity is to have a comprehensive policy in place that covers appropriate behavior. After that, it's a matter of executing the proper steps per the plan when something happens.

Situations involving prohibited content or activities are not easy to address. The accused person might get angry or confrontational, so it's important always to have the right people there to help manage and defuse the situation. If you feel that the situation is severe enough to worry about your own personal safety, don't be afraid to involve the police. While the situation needs to be handled, there's no sense in putting yourself in direct danger to do so.

Creating a Prohibited Content Policy

Creating a policy is the most important part of dealing with prohibited content or actions. Without a policy in place that specifically defines what is and what isn't allowed, and what actions will be taken when a violation of the policy occurs, you don't really have a leg to stand on when a situation happens.

What is contained in the policy depends on the organization for which you work. Generally speaking, if something violates an existing federal or local law, it probably isn't appropriate for your network either. Many organizations also have strict policies against the possession of pornographic or hate-related materials on the organization's property. Some go further than that, banning personal files such as downloaded music or movies on work computers. Regardless of what is on your policy, always ensure that you have buy-in from very senior management so that the policy will be considered valid.

Here are some specific examples of content that might be prohibited:

- Adult content
- Content that advocates violence against an individual, group, or organization
- Unlicensed copyrighted material
- Content related to drugs, alcohol, tobacco, or gambling
- Content about hacking, cracking, or other illegal computer activity
- Violent or weapons-related content

A good policy will also contain the action steps to be taken if prohibited content or activity is spotted. For example, what should you do if you find porn on someone's work laptop?

The policy should explicitly outline the punishment for performing specific actions or possessing specific content. The appropriate penalty may very well be based on the type of content found. Something that is deemed mildly offensive might result in a verbal or written warning for the first offense and a more severe sentence for the second offense. If your company has a zero-tolerance policy, then employees may be terminated and possibly subject to legal action.

Finally, after the policy has been established, it's critical to ensure that all employees are aware of it and have proper training. In fact, it's highly recommended that you have all employees sign a disclosure saying they have read and understand the policy, and that the signed document be kept in their human resources file. Many organizations also require that employees review the policy yearly and re-sign the affidavit as well.

Incident Response

If you have your policy in place, then your *incident response* plan should be relatively scripted. It might not be easy to deal with, but the steps you should take should be outlined for you. Professionalism should be maintained during the incident. This is a good time to remind you that people will be looking at your reaction as well as your actions, and professionalism will define how you dealt with the incident. If you see prohibited content and start giggling and walk away, that probably doesn't reflect well on you. Always remember that others are watching you. The specific steps that you take will depend on your policy. The following sections describe the best practices for dealing with security incidents.

Detecting the Incident

An incident can be detected in several different ways. The preceding section used the example of pornographic content on a laptop or mobile device. This is a great illustration of an example of passive detection. You were not looking for this material, but you found it, and now you must respond. A full list of incident detections is detailed in the following:

Passive *Passive detection* is used to clue in on what happened and how it happened. Passive detection can be when you are reviewing logs, content, or events and discover an incident. Another example of passive detection in a physical security context is the use of a video camera. The video camera will show you an event in the past, but it is normally not actively looking for an event. An example of passive detection in a digital context is the use of file hashes. If a file hash does not match the original, then someone or something has tampered with the original content.

Active *Active detection* is used to actively look for security incidents. This type of detection in a physical security context is similar to having a security guard patrol your premises. A digital example might be the use of an *intrusion detection system (IDS)*

running on your firewall that alerts you to an attempted break-in. This type of detection is used to catch a perpetrator in the act of a security incident.

Proactive *Proactive detection* is the act of proactively securing possible holes in security. *Penetration testing*, also known as *pentesting*, involves someone trying to penetrate either the physical or software defenses of the organization. The exercise is done to strengthen the organization's security for either the physical or digital defenses. Proactive detection normally does not create a security incident, unless the potential for the loss of information was discovered during the process.

Responding to the Incident

Once an incident is detected using the methods mentioned in the preceding section (passive, active, or proactive), or it's detected through dumb luck, it's time to spring into action and respond to the incident. The person responding to the incident, called the *first responder*, should be versed in how to collect evidence in the order of volatility. If the evidence is not collected, for example, from a computer's RAM, and the computer reboots, the evidence will be gone. The *order of volatility* is as follows:

Memory Contents Evidence of the incident located in the RAM of the computer or system. This evidence can be a file that has not been saved to the filesystem. It can also be a process running in RAM that requires a memory dump for further analysis.

Swap Files/Virtual Memory Evidence that has been saved out to a swap file/virtual memory, because the process has been backgrounded or is no longer in foreground RAM. The swap file/paging file should be copied to removable media.

Network Processes Evidence that is part of a network process, such as a browser redirection or spam or active network communications. This is where most of the incidents occur on a network. This information should be saved to removable media.

System Processes Evidence that is part of a system, such as an exploit that has been rooted into the operating system. This is also a common place for incidents to occur. The information should be separated out of the system and saved to removable media.

Filesystem Information This evidence is in the form of files on a hard drive. An example of filesystem information is illegal content that has been saved to the filesystem. This information must be preserved in its original form but also copied for evidence purposes.

Raw Disk Blocks This evidence is in the form of a block-level copy of the data. Once an incident is discovered and all the other volatile evidence has been collected, the last step is to create an image of the affected system. This can be done with a tool called a *write-blocker* (which makes the media read-only) and specialized software that copies data at the block level.

The preceding evidence can be collected with the sophisticated tools that only a highly trained first responder might have on hand. However, not all the tools need to be complex tools. A simple camera can preserve information. For example, if you walk up to a system that displays a ransomware screen, your first reaction should be to take a photo of the screen. If you press a key, the ransomware could crash and disappear. Photo evidence of data and processes that are loaded in RAM is a good alternative to not having any evidence, but having a digital copy of the evidence is preferred.

The act of photographing the scene should not be limited to just the computer screen; anything relevant to the incident should be photographed as evidence. Time and date stamps should be overlaid onto the image. This is normally a function of any camera. However, if you can't digitally record the time stamps, a simple alternative is to include a watch in the frame of the photo.

You should take notes with a pad and pen, recording the initial scene, including time and date. Create a chronology of the discovery and collection of the evidence. Remember, any of this could potentially be used in a court of law. The underlying premise is to record as much evidence as possible before the crime scene is tainted by others.

Following the Right Chain of Custody

The removed materials should be secured and turned over to the proper authorities. Depending on the situation, materials may be held in a safe, locked location at the office, or they may need to be turned over to local authorities. Have a documented procedure in place to follow, given a situation.

The materials that are deemed as evidence should be well documented as to why they are considered evidence. The *chain of custody* documentation should define the following:

- Who obtained the evidence
- Who secured the evidence
- Who controlled the evidence during the entire process

If the evidence is moved, the chain of custody documentation should reflect the following:

- The reasons it was moved
- Who moved it
- How it was secured
- Who controlled it

The chain of custody must be maintained at all times. If a chain of custody of the evidence is not maintained, the evidence may not be admissible in a court of law.

Reporting the Incident

Once you've collected the initial set of evidence and in the order of volatility, it's time to report the incident. The incident should be reported to management and a decision should be made whether to involve law enforcement. Involving law enforcement is, of course, dependent on the severity of the incident. A piece of malware that has infected a

single machine is a bit different than malware that has infected an entire network. Your management might elect not to involve law enforcement at all.

Regardless of the direction management takes, the evidence collected, as well as notes taken during the collection, will help an escalation team or law enforcement to proceed in building a case. The goal of the first responder is to collect evidence that answers the following questions:

- What happened?
- When did it happen?
- How did it happen?
- Who made it happen?
- Why did it happen?

An escalation team or law enforcement's job is to fill in the blanks by using the evidence. If the evidence is complete and concise, it will be used to build a case against the threat agent. The ultimate goal is to stop a future incident from happening.

Recovering from the Incident

Your most important task is to recover from the incident. If your critical ordering system was affected during the incident, it's your job to get it back online. If the incident involved one computer that is used by a task worker, then it's your job to get it back up and running. You might notice a common theme here: it's your job to get things back to normal after the incident. Once you can get the flow of information flowing again, you can move on to remediating the incident.

During the recovery from the incident, you may have to make changes to the network or systems that support the clients. Any changes during the recovery process should be documented thoroughly. Documenting these changes is important if you are submitting claims to insurance, assessing damages, or looking for future reparations.

Remediating the Incident

All components affected by the incident should be remediated to ensure that all traces of the incident have been removed. Remediation can be as simple as adding firewall rules to the firewall, or it can involve formatting a server and reloading it. The steps to remediation should also include steps to prevent the incident from happening in the future.

Documenting the Incident

Before, during, and after the incident, the documentation process should begin. You should collect as much information as possible, as a formal incident report will eventually be formulated defining the following key elements.

- Date and time
- Summary
- Root cause
- Actions taken

- Remediation
- Services impacted
- Recommendations

It really doesn't matter how you collect information for documentation purposes. It can be pad and pen or something more elaborate. The only stipulation is that the documentation should not be on a system that can be affected by the incident. An offline laptop is fine, as long as the laptop is never introduced to the network affected by the ongoing incident. This could jeopardize all of the documentation efforts and hinder the outcome.

Reviewing the Incident

The final step for incident response is to review all the documentation and findings of the incident—a process often called a hot-wash meeting. During a hot-wash meeting, the incident response team should talk about what has been done properly during the incident and what procedures should be changed for future incidents. These meetings should be constructive and support standards of excellence for the incident response team.

Another key goal of the review process is to identify threats similar to the characteristics of the incident. If an employee entered credentials into a phishing page, what measures are in place to prevent this from happening to others in the organization? You may have rules in place for this particular phishing email, but are your employees trained for future incidents similar to this? If not, end-user training may be required.

Best Practices for Incident Response

Now that you have a good understanding of the process involved with incident response, let's look at some of the best practices. The following best practices should be applied to all elements of the incident response process:

Follow your policies exactly as they are written. Yes, we've already said this several times. It's crucial that you do this. Not following the policies and procedures can derail your case against the offender and possibly set you up for problems as well.

If you are the first responder, get a verifier. Your first priority as the first responder is to identify the improper activity or content. Then you should always get someone else to verify the material or action so that it doesn't turn into a situation of your word against someone else's. Report the situation immediately through the proper channels.

Preserve the data or device. The data or device should immediately be removed from the possession of the offending party and preserved. This will ensure that the data doesn't mysteriously disappear before the proper parties are notified.

Use documentation. Document everything that could be relevant to the situation. Many companies have standard documentation that is used in incident response in order to be sure that the responder captures important information and does not forget to ask critical questions or look for vital clues.



Always be sure to follow government regulations as well as corporate end-user policies and security best practices when dealing with sensitive information.

Common Agreements

There are several agreements that you will need to be familiar with for the CompTIA exam. This short list does not constitute all of the agreements you might have to abide by as a network professional. These are just the most common that you might need to follow or create.

Nondisclosure Agreement (NDA)

When an organization hires a new employee, the employee may be privilege to certain proprietary information. This information could be a company's secret sauce, passwords to administrative systems, or customer information. A nondisclosure agreement (NDA) legally protects an organization from an employee disclosing this information to a person outside the organization. The NDA is a document that defines the consequences of disclosure by termination and legal recourse. The NDA should be developed by the organization's legal counsel, HR department, and stakeholders of proprietary information in the organization. These documents are typically signed by the employee during the onboarding process.

Mutual Nondisclosure Agreement (MNDA)

A mutual nondisclosure agreement (MNDA) is a legal agreement in which both parties agree to protect each other's confidential information. This type of agreement is often used in situations where both parties need to share sensitive information with one another. This is commonly found with business partnership, joint venture, or collaboration. The mutual NDA ensures that both parties have a legal obligation to keep the shared information confidential and not disclose it to a third party.

An MNDA differs from a standard NDA, because an NDA typically involves only one party disclosing confidential information to the other. In this case, the receiving party agrees to keep the information secret and not use it for any unauthorized purposes. While a standard NDA imposes confidentiality obligations on just one party, an MNDA imposes these obligations on both parties equally. The key difference between the two types of NDAs lies in the reciprocity of the mutual NDA.

Acceptable Use Policy (AUP)

There can be many different resources employees use that the organization supplies to conduct its business. An *acceptable use policy (AUP)* helps protect the organization from employee abuse of these resources. Email is an example of one such resource. It is generally not acceptable for an employee to use an organization's email for religious, political, or personal causes; illegal activities; or commercial use outside of the organization's interest.

The AUP should be developed by the organization's legal counsel, HR department, and IT department. The systems that should be included in the AUP are telephone, Internet, email, and subscription services. The AUP might not be exclusive to electronic resources; the organization might also include postage and other nonelectronic resources that could be abused. The AUP is typically signed or electronically acknowledged during the employee's onboarding process, or when a resource is accessed that could be misused.

Regulations and Compliance

Regulations are rules imposed on your organization by an outside agency, like a certifying board or a government entity, and they're usually totally rigid and immutable. The list of possible regulations that your organization could be subjected to is so exhaustively long, there's no way I can include them all in this book. Different regulations exist for different types of organizations, depending on whether they're corporate, nonprofit, scientific, educational, legal, governmental, and so on, and they also vary by where the organization is located.

For instance, U.S. governmental regulations vary by county and state, federal regulations are piled on top of those, and many other countries have multiple regulatory bodies as well. The Sarbanes–Oxley Act of 2002 (SOX) is an example of a regulation system imposed on all publicly traded companies in the United States. Its main goal was to ensure corporate responsibility and sound accounting practices, and although that may not sound like it would have much of an effect on your IT department, it does because a lot of the provisions in this act target the retention and protection of data. Believe me, something as innocent sounding as deleting old emails could get you in trouble—if any of them could've remotely had a material impact on the organization's financial disclosures, deleting them could actually be breaking the law. All good to know, so be aware, and be careful!

The penalties of noncompliance to regulations can be in the form of financial penalties, reputation damage, loss of services, loss of licensing, and ultimately disruptions to operations. Therefore, it is paramount to the continued operation of our organizations that we comply with safeguarding the data entrusted.

Splash Screens

Splash screens are used by organizations to communicate compliance-related information. These screens can be customized to display various types of content, such as legal disclaimers, privacy policies, or terms of service (ToS), which users must acknowledge before accessing the system. This ensures that users are consistently reminded of their obligations and the organization's policies. The organization will typically use these splash screens on a captive portal, before an employee or guest is granted access to a network. The splash screen can also include an AUP reminding the end-user of the acceptable use.

For example, a healthcare organization that needs to comply with healthcare regulations, might implement a splash screen that appears each time a user logs into their system. This will help remind them of the importance of protecting patient information and adhering to data privacy standards. The splash screen could include a brief summary of key compliance points and a prompt for the user to acknowledge their understanding before proceeding.

This not only reinforces the organization's compliance policies, but also helps to create a culture of accountability and awareness among the organization's employees.

Healthcare organizations are not the only organizations that can benefit from splash screens. This approach can be particularly effective in industries with stringent regulatory requirements, such as finance, healthcare, and legal sectors, just to name a few.

Managing Software Licenses

When you buy an application, you aren't actually buying the application. Instead, you're buying the right to use the application in a limited way, as prescribed by its licensing agreement. Most people don't read these licensing agreements closely, but suffice it to say, they're pretty slanted in favor of the software manufacturer.

Don't like the terms? Too bad. No negotiation is allowed. If you don't accept the *end-user license agreement (EULA)*, your only recourse is to return the software for a refund. (Most vendors will refuse to take back an opened box. Still, the software manufacturer is required to take it back and refund your money if you reject the licensing. This is true of programs purchased online as well.)



Many companies rely on *digital rights management (DRM)* to protect digital assets, such as online photos or videos. DRM is not as established as licensing agreements are, but you should still respect the property of the owners of digital content.

Although the majority of the applications that you acquire will probably be commercial products, there are a number of alternatives to commercial software sales. Here are some of the license types that you may encounter:

Freeware *Freeware* is software that is completely free. On a small scale, you can get such software from download sites such as <https://download.cnet.com> or from the creator's personal website. Large companies like Google and Microsoft also sometimes offer products for free, because it serves the company's interests to have a lot of people using their software. Examples include Google Chrome and Microsoft Internet Explorer. Freeware doesn't include source code, and users aren't allowed to modify the application.

Open Source Open source software is freer than free: not only is the application free, but the source code (code used by programmers) is also shared to encourage others to contribute to the future development and improvement of the application. OSs such as Linux and applications such as OpenOffice fit this category. Open source software can't be sold, although it can be bundled with commercial products that are sold.

Shareware Shareware is software that provides a free trial, with the expectation that you'll pay for it if you like it and decide to keep it. In some cases, a shareware version isn't the full product; in other cases, it expires after a certain amount of time. Some shareware provides a full and unlimited version, with payment requested on the honor system.

Multiuser This is commercial software that you're allowed to install on more than one computer. For example, some versions of Microsoft Office allow you to install the same copy on two or three PCs.

Single User This is commercial software for which the license restricts installation to a single PC. A common misconception is that a single-user license allows you to install the software on more than one computer as long as you use only one instance at a time, but that's not accurate. Commercial products sometimes have activation systems that lock the software to a specific PC once it's installed, so you can't install it elsewhere. Microsoft Office is a good example of commercial software.

Concurrent This license allows the software to be installed on many PCs but used concurrently by a smaller number. For example, you may have 1,000 computers with the application installed, but only 100 users can use it simultaneously. This is useful in situations in which everyone needs to have an application available but the application gets very little actual use.


Perpetual A perpetual or non-expiring license is a license that can be purchased for a software product. Many video games still use this license model. You purchase the software and the right to use the software for life. Software vendors that use perpetual licensing will also offer software upgrades or patches for the first year, as customary. However, any upgrades after the first year are generally covered under a support or maintenance fee.

Corporate, Campus, or Site These are enterprise licenses, which permit an organization to install the application on an agreed-upon number of PCs. For example, a school may buy a site license of an antivirus program and allow all students to download and install it freely to ensure that the school's network remains virus free. Microsoft calls this a volume license key.

Valid License A valid license is any software license that is current in its term and not counterfeit. As long as the license is not used already, it is valid for activation for the software product it is associated with.

If you buy any sort of commercial software, you will receive a *product key*, which you will need to enter during installation or the first time the application is opened. The product key might be emailed to you, or it could be located on the physical media if you got an installation CD-ROM, DVD, or thumb drive. Figure 9.12 shows an example of a product key.

In a corporate environment, license management is a critical responsibility. The company may spend thousands or even millions of dollars on software licenses. Money could be wasted on unused licenses, or if the company's computers have unlicensed software, it could result in huge fines. Ignorance is not a legal excuse in this area.

FIGURE 9.12 A Microsoft product key

To avoid these problems, it may be best for your company to purchase a software asset management tool, such as Microsoft's Software Asset Management (SAM) guide (www.microsoft.com/en-us/download/details.aspx?id=31382), License Manager by License Dashboard (www.licensedashboard.com), or FlexNet Manager by Flexera Software (www.flexera.com/products/flexnet-manager). In general, here are the steps to take for proper license management:

1. Build a database of all licenses owned by your company. This includes what type of license it is, license numbers, and expiration dates, if any.
2. Perform an inventory of all licensed software installed on your computers. (This can be quite an effort.)
3. Compare the license list to the installation base.
4. For any gaps where you have unlicensed software, either remove the software or procure licenses.

Because of the potential for heavy fines, many companies prohibit the installation of software on client computers unless specifically authorized by a manager or the IT department.



Real World Scenario

How Do I Buy the Right Licenses?

Consider this situation: After talking to your boss about software licenses, you decide to investigate the office productivity software on your department's computers. You are unable to find proper documentation that the correct software licenses were purchased, leading you to wonder whether the company has the right licenses. Your boss wants to avoid any

potential legal issues and asks you to go buy enough copies of the latest version of Microsoft Office for the 20 users in your department. What do you do?

The first question to ask yourself is, do all 20 people use their computers at the same time? If not, then you might be able to purchase a concurrent license for fewer than 20 users. If there's the possibility of all 20 users needing Office at once, then you'll definitely need licenses for everyone. You might purchase a corporate license as well. Now that you have figured out what to buy, how do you do it?

One option is to go to the local computer store, load up 20 boxes of Office (if they have that many in stock), and trudge up to the cashier. A second option is to go to Microsoft's volume licensing site, at www.microsoft.com/en-us/licensing, to learn about purchasing multiple licenses. Microsoft will direct you to an authorized reseller so that you can purchase the licenses. The authorized reseller will register the list of license numbers with Microsoft's licensing center. Now you need just one physical (or downloaded) copy to perform the installation.

Managing Sensitive Information

As an IT manager, you will very likely have access to information that you will need to keep closely guarded. For example, you might have access to username and/or password lists, medical or educational records, addresses and phone numbers, or employee records. It's your responsibility to ensure that sensitive information does not get released into the wrong hands. On the flip side, you may encounter information that's sensitive because it's prohibited or illegal. You need to know how to react in those situations as well.

Much of the data you come into contact with could be regulated outside of the organization's internal policies. Regulated data must be identified as it enters your network, and the proper operating procedures should be followed. The operating procedures for such data should be constructed so that you can adhere to the data's regulatory compliance rules. The regulatory rules for compliance can be at any local, state, or federal level. Many of the regulations in the following sections are at a federal level of compliance.

Personally Identifiable Information

Personally identifiable information (PII) is anything that can be used to identify an individual person on its own or in context with other information. This includes someone's name, address, other contact information; the names of family members; and other details that people would consider private.

PII should always be kept confidential and secure. It seems like every few months or so we see news stories of data breaches at big companies resulting in stolen credit card data or username and contact lists. This information finds its way into hackers' hands and causes millions of people grief and monetary damages. Be sure that this information is properly secured and can be accessed only by authorized personnel.

Personal Government-Issued Information

Any personal information contained in a document issued by a government or state is considered personal government-issued information. Examples of government-issued documents are defined as a birth certificate, Social Security card, identification card, driver's license, resident card, taxpayer ID number, or password, just to name a few. The category of information is very broad and can overlap with other types of protected data, such as protected health information or other PII.

The threat to personal government-issued information being compromised is that the information is how a person is defined by the government. The theft of a Social Security number is a direct theft of identity in the eyes of the government. Tax records to credit information is tied to a Social Security number. Therefore, like any PII, this data should also be kept confidential and secure.

Payment Card Industry Data Security Standard

Payment Card Industry Data Security Standard (PCI DSS) is a standard of processes and procedures used to handle data related to transactions using payment cards. A payment card is any card that allows the transfer of money for goods or services. Types of payment cards include credit cards, debit cards, or even store gift cards.

PCI DSS compliance is not enforced by government entities. PCI DSS compliance is actually enforced by banks and creditors. Merchants must comply with the PCI DSS standard to maintain payment card services. If a merchant does not comply with PCI DSS standards and a breach occurs, the merchant can be fined by the banks. Once a breach of PCI data occurs, then local, state, and federal laws can apply to the merchant. For example, some laws require the merchant to pay for credit-monitoring services for victims after a breach.

General Data Protection Regulation

The *General Data Protection Regulation (GDPR)* is a European Union (EU) law governing how consumer data can be used and protected. The GDPR was created primarily to protect citizens of the European Union. It applies to anyone involved in the processing of data based on the citizens of the European Union, regardless of where the organization is located.

The GDPR recommends that organizations hire a *data protection officer (DPO)*. This person is the point of contact for all compliance with GDPR, as well as any other compliances your organization falls under. The underlying goal is to achieve consent from the end user of your product or service. Consent to collect information must be proven by an organization beyond a shadow of doubt. This means that if someone visits your website from the European Union, you must receive consent in clear language to even place a cookie in their web browser. The DPO is responsible for coordinating this language, as well as the life cycle of any data that is collected.

Protected Health Information

Protected health information (PHI), also known as *personal health information*, refers to any information used in the health care industry to describe a patient or ailment. This information can be considered “the patient chart” you always see on television. However, *electronic health records (EHRs)* go way beyond the current condition of a patient; they describe a person from the cradle to the grave.

Electronic health records (EHRs) are used to record a patient’s vitals every time the patient visits a doctor’s office. They represent historical information about patients, as well as billing information used by health care providers. This makes the EHR extremely valuable to a hacker and represents a large makeup of identity theft.

This type of identity theft is really dangerous! Your diagnosis could be determined based upon vitals, allergies, or conditions that are recorded from a person who assumes your identity. Try to explain to the insurance company that your gallbladder needs to be removed, but their records show they paid to have it taken out already.

PHI can also be used to track the statistics of larger groups of people. However, the data must be anonymized first, before it is put into a publicly addressable database for these statistics and studies. Information such as names, Social Security numbers, phone numbers, health insurance information, account numbers, and specific geographical information must be removed. Interesting enough, geographical information can include only the first three digits of a zip code.

Data Retention Requirements

As you’ve learned, there are all different forms of data and they are regulated for sensitivity and confidentiality purposes. The storage of data is also regulated depending on the type of data. The data to be retained is usually transactional in nature, such as a financial transaction, but it can also be data that has touched your system.

Regulations are not the only governing factor for data retention—your company can also have an internal requirement. An example of an internal retention requirement might be for memorandums, access control records, or even video footage. There are different reasons for adhering to an internal data retention requirement, but they all shield the organization from liability in some way. This is very apparent in legal situations. If you keep all data indefinitely on old backups and there is a lawsuit, you will be liable to restore records in a timely fashion, even if those old systems are no longer supported. If you don’t produce the data in a timely fashion, you could default a lawsuit in favor of the plaintiff. These discoveries of information are considered e-discovery by court proceedings.

As an IT professional, it is your responsibility to work with legal counsel in your organization to define data retention periods. However, before you can define the data retention periods, you will need to identify the data to be protected. A document profile should be created that clearly defines the types of data you deal with on a daily basis. This stage of the process is where you can also identify outside governing requirements for retention. Once the document profile is created, you should start tagging data in your environment with the

data types. This may be as simple as naming the email backup job email data. You should then set up a hard policy to control how long the data is kept. Keep in mind that holding the data too long is just as bad as not holding the data long enough.

Summary

This chapter covered three areas of operational procedures that you should integrate into your work:

- Safety procedures
- Environmental impacts and controls
- Licensing, policies, and sensitive materials

First, we looked at the importance of safety procedures. Safety is about protecting you from harm as well as protecting your computer components from getting damaged. Then, we outlined some methods to apply the policies and procedures for a safe working environment, and identified potential safety hazards. Included were preventing *electrostatic discharge (ESD)* and *electromagnetic interference (EMI)*, creating a safe work environment, and properly handling computer equipment.

Safety involves you and your coworkers, but it also includes environmental issues. The environment can have a harmful effect on computers, but computers can also greatly harm the environment. You need to be familiar with the importance of material safety data sheets (MSDSs) as well as the proper disposal procedures for batteries, display devices, and chemical solvents and cans. These items need to be kept out of the environment because of the damage that they can cause.

Finally, we looked at potential legal issues. Failure to follow certain procedures can expose you or your company to legal proceedings. Make sure that all the software on your computers is legal and licensed and that the computers contain no illegal or prohibited materials. You may also need to protect personally identifiable information, depending on the type of data you have. When incidents happen, you need to know how to respond properly to mitigate the issue.

Exam Essentials

Know which computer components are particularly dangerous to technicians. The most dangerous components are the power supply and the monitor. Both are capable of storing lethal charges of electricity, even when unplugged. You also need to be aware of parts that get incredibly hot, such as the processor, which can cause severe burns if touched.

Understand where to find safety information regarding chemicals. You can find this information in a material safety data sheet (MSDS). An MSDS might not have come with your purchase, but most suppliers will gladly supply one.

Know which tool to use for which job. The majority of computer repair jobs can be handled with nothing more than a Phillips-head screwdriver. However, you might need cutters, extra light, or a mirror for some jobs. Avoid using magnetic-tipped tools.

Understand methods to help prevent ESD. One of the biggest and most common dangers of electronic components is electrostatic discharge (ESD). You can employ several methods to help avoid ESD-related problems, such as grounding yourself; using an antistatic wrist strap, bag, or mat; and controlling the humidity levels.

Know the proper disposal procedures for used computer parts, batteries, and chemical solvents. The specific disposal procedure depends on the item. However, the safe answer is always to recycle the component, not to throw it in the trash.

Know the differences between license types. Open source applications don't require a purchased license, whereas commercial applications do require a purchased license. Some licenses are personal licenses, meaning the applications are for use by one person on one computer, whereas others are enterprise licenses, meaning the applications can be used on multiple computers at the same time.

Understand how to handle prohibited content or activity. Always have policies and procedures in place to deal with prohibited content or activity. When an incident happens, follow the procedures, report through proper channels, preserve the data or device using the order or volatility, and follow the chain of custody.

Know the various types of sensitive information. You should know that personal identifiable information (PII) is any information used to identify an individual. The Payment Card Industry Data Security Standard (PCI DSS) is a standard of policies and procedures used by the payment industry. General Data Protection Regulation (GDPR) is a European Union (EU) law governing how EU citizen data is used. Protected health information (PHI) is any information used to describe a patient or a patient's ailment.

Know the common business agreements. These agreements include nondisclosure agreements (NDA), mutual nondisclosure agreements (MNDA), and acceptable use policy (AUP). You should understand how they are used to protect organizations and individuals.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. You have dropped a screw into a tight location of a computer. How should you retrieve it?
 - A. Use a magnetic-tipped screwdriver.
 - B. Use a magnetic grabber.
 - C. Use a three-pronged grabber.
 - D. Shake the computer until it falls out.
2. Your company's AUP states that employees must not use company resources for personal gain. You have a side business and want to use your work computer to manage it. What should you do?
 - A. Use your work computer to manage your side business.
 - B. Use your personal computer to manage your side business.
 - C. Ask your manager for permission to use your work computer.
 - D. Ignore the AUP and use your work computer anyway.
3. One of your coworkers just spilled a chemical solvent in a warehouse, and you have been asked to help clean it up. Which of the following must contain information about a chemical solvent's emergency cleanup procedures?
 - A. OSHA
 - B. MSDS
 - C. Product label
 - D. CRT
4. You are purchasing an inkjet printer cartridge for home use that you know has an MSDS. How do you obtain the MSDS for this product?
 - A. The store is required to give you one at the time of purchase.
 - B. It's contained in the packaging of the printer cartridge.
 - C. You are not legally allowed to have an MSDS for this product.
 - D. You should visit the printer cartridge manufacturer's website.
5. In the interest of a safe work environment, which of the following should you report? (Choose two.)
 - A. An accident
 - B. A near-accident
 - C. Spills on the floor inside a building
 - D. Rain forecasted for a workday

6. What is the approximate minimum level of static charge for humans to feel a shock?
 - A. 300 volts
 - B. 3,000 volts
 - C. 30,000 volts
 - D. 300,000 volts
7. Your work environment has been unusually dry lately, and several components have been damaged by ESD. Your team has been asked to be extra careful about ESD damage. Which of the following measures can be implemented to reduce the risk of ESD? (Choose two.)
 - A. Use an antistatic wrist strap.
 - B. Use an antistatic bag.
 - C. Spray disinfectant spray.
 - D. Shuffle your feet.
8. Which of the following are OSHA requirements for a safe work environment that must be followed by employers? (Choose two.)
 - A. Attend yearly OSHA safe work environment seminars.
 - B. Provide properly maintained tools and equipment.
 - C. Have an OSHA employee stationed within 5 miles of the facility.
 - D. Display an OSHA poster in a prominent location.
9. Your office just added 20 new workstations, and your manager has put you in charge of configuring them. The users need to have Microsoft Office installed. What should you do to install Microsoft Office properly on these computers?
 - A. Ensure that the company has the proper licenses to install 20 additional copies.
 - B. Agree with the open source license agreement during installation.
 - C. Use the personal license key from an existing system to install Office on the new computers.
 - D. Follow normal installation procedures; nothing else needs to be done.
10. Your office is moving from one floor of a building to another, and you are part of the moving crew. When moving computer equipment, which of the following are good procedures to follow? (Choose two.)
 - A. Lift by bending over at the waist.
 - B. Carry CRT monitors with the glass face away from your body.
 - C. Use a cart for heavy objects.
 - D. Ensure that there are no safety hazards in your path.
11. You just removed four AA alkaline batteries from a remote-control device. What is the recommended way to dispose of these batteries?
 - A. Throw them in the trash.
 - B. Incinerate them.
 - C. Take them to a recycling center.
 - D. Flush them down the toilet.

12. When replacing a hard drive, you discover prohibited material on a user's laptop. What should you do first? (Choose two.)
- A. Destroy the prohibited material.
 - B. Confiscate and preserve the prohibited material.
 - C. Confront the user about the material.
 - D. Report the prohibited material through the proper channels.
13. You need to investigate how to protect credit card data on your network. Which information should you research?
- A. PCI DSS
 - B. GDPR
 - C. PHI
 - D. PII
14. Which class of fire extinguisher is recommended for use in a wood and paper fire?
- A. A
 - B. B
 - C. C
 - D. D
15. Which of the following are common types of screwdrivers? (Choose two.)
- A. Circular
 - B. Phillips
 - C. Torx
 - D. Helix
16. Which of the following are elements of a good workplace safety plan? (Choose two.)
- A. Periodic workplace inspections
 - B. A safety and health training program
 - C. Retribution for employees who report violations
 - D. An independent third-party auditor of the safety plan
17. What is the recommended use policy on magnetic-tipped screwdrivers inside computers?
- A. Do not use them.
 - B. It's okay to use them, but keep them away from the processor.
 - C. It's okay to use them, but keep them away from the RAM.
 - D. It's okay to use them, but only if they're of the powered variety.
18. Which of the following are usually contained on an MSDS? (Choose two.)
- A. Freezing point
 - B. Handling and storage instructions
 - C. Personal protection instructions
 - D. Weight

- 19.** Which of the following are OSHA requirements for a safe work environment that must be followed by employees? (Choose two.)
- A.** Immediately report all accidents to OSHA.
 - B.** Use protective gear and equipment.
 - C.** Attend safety training.
 - D.** Follow all employer-implemented health and safety rules.
- 20.** Which of the following types of batteries are not considered environmental hazards?
- A.** Alkaline
 - B.** Nickel-metal hydride (NiMH)
 - C.** Nickel-cadmium (NiCd)
 - D.** Button cell

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

One of your office coworkers recently tripped on a power cord and injured himself. What should you do to find potential trip hazards in your office? Once the hazards are identified, what actions should you take?



Chapter 10

Documentation and Professionalism

THE FOLLOWING COMPTIA A+ 220-1202 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **4.0 Operational Procedures**

✓ **4.1 Given a scenario, implement best practices associated with documentation and support systems information management.**

- Ticketing systems
 - User information
 - Device information
 - Description of issues
 - Categories
 - Severity
 - Escalation levels
 - Clear, concise written communication
 - Issue description
 - Progress notes
 - Issue resolution
- Asset management
 - Inventory lists
 - Configuration management database (CMDB)
 - Asset tags and IDs
 - Procurement life cycle
 - Warranty and licensing
 - Assigned users
- Types of documents
 - Incident reports
 - Standard operating procedures (SOPs)



- Software package custom installation procedure
- New user/onboarding setup checklist
- User off-boarding checklist
- Service-level agreements (SLAs)
 - Internal
 - External/third-party
- Knowledge base/articles

✓ **4.2 Given a scenario, apply change management procedures.**

- Documented business processes
 - Rollback plan
 - Backup plan
 - Sandbox testing
 - Responsible staff members
- Change management
 - Request forms
 - Purpose of the change
 - Scope of the change
 - Change type
 - Standard change
 - Normal change
 - Emergency change
 - Date and time of change
 - Change freeze
 - Maintenance windows
 - Affected systems/impact
 - Risk analysis
 - Risk level
 - Change board approvals
 - Implementation



- Peer review
- End-user acceptance

✓ **4.3 Given a scenario, implement workstation backup and recovery methods.**

- Backup
 - Full
 - Incremental
 - Differential
 - Synthetic full
- Recovery
 - In-place/overwrite
 - Alternative location
- Backup testing
 - Frequency
- Backup rotation schemes
 - On site vs. off site
 - Grandfather-father-son (GFS)
 - 3-2-1 backup rule

✓ **4.7 Given a scenario, use proper communication techniques and professionalism.**

- Present a professional appearance and wear appropriate attire.
 - Match the required attire of the given environment
 - Formal
 - Business casual
- Use proper language and avoid jargon, acronyms, and slang, when applicable
- Maintain a positive attitude/project confidence
- Actively listen, take notes, and avoid interrupting the customer
- Be culturally sensitive
 - Use appropriate professional titles, when applicable
- Be on time (if late, contact the customer)



- Avoid distractions
 - Personal calls
 - Texting/social media sites
 - Personal interruptions
- Appropriately deal with difficult customers or situations
 - Do not argue with customers and/or be defensive
 - Avoid dismissing customer problems
 - Avoid being judgmental
 - Clarify customer statements (i.e., ask open-ended questions to narrow the scope of the problem, restate the issue, or question to verify understanding)
 - Use discretion and professionalism when discussing experiences/encounters.
- Set and meet expectations/timeline and communicate status with the customer
 - Offer repair/replacement options, as needed
 - Provide proper documentation on the services provided
 - Follow up with customer/user at a later date to verify satisfaction
- Appropriately handle customers' confidential and private materials
 - Located on a computer, desktop, printer, etc.



Every day at work, we do what it takes to get the job done. As IT professionals, we have millions of facts crammed into our heads about how various hardware components work and what software configuration settings work best for our systems. We know which servers or workstations give us constant trouble and which end users need more help than others. We are counted on to be experts in knowing *what* to do to keep computers and networks running smoothly. And even though we don't tend to think about it overtly, we need to be experts in *how* we get things done as well. Even though the *how* might not be at the top of your mind every day (hopefully you don't go to work every day thinking "Okay, don't get killed by a monitor" or "Let's see if I can be nice to someone today"), but it should be integrated into your work processes. Operational procedures define the *how*, and they provide guidance on the proper ways to get various tasks accomplished.

In this chapter, we will describe documentation used to document the network and the policies that you will encounter and need to enforce. We will then look at change management practices used to make changes to the network that will affect other business units in your organization.

We will also look at disaster prevention and recovery. In the process, we will cover topics such as backup and recovery, power resiliency, and cloud storage versus local storage.

In the final part of this chapter, we'll switch to discussing professionalism and communication and focus on topics that you need to know for your exam study. Applying the skills learned here will help you pass the exam, but on a more practical level, it will help you become a better technician and possibly further advance your career.



Exam Objective 4.1, "Given a scenario, implement best practices associated with documentation and support systems information management," requires you to know ticketing systems and the various elements they contain. You must know the various elements of asset management and the types of documents you will use daily.

Documentation and Support

Creating documentation of the network and systems you work on is the last step to many projects. However, it is the most overlooked, or it is considered secondary to fixing new problems or upgrading to the next system. The documentation produced has positive effects

on production and problem solving. It is used so that you don't need to redo the step of discovery the next time you have a problem. Documentation also allows you to have others work on a problem with the same distinctive view you had at the time of documentation.

Some of the documentation you create will help to develop policies and procedures that others will need to follow. Throughout this book, we have discussed the hard controls of these policies. For example, when you implement a password policy, you can dictate that a password be complex and of a certain length. A written policy is a soft control that might detail how to create a complex password. In this chapter, we will look at several different policies that you will come across as a technician.

Ticketing Systems

One of the most important functions of the information technology (IT) department is to solve problems and fulfill end-user requests. The core of this function is the ability to receive incoming requests, track progress, and assure that problems and requests are solved and completed. It is common to find a ticketing system at the heart of this function called help desk services or support services. A ticketing system adds other benefits to the function of the help desk, such as accountability, reporting, collaboration, and escalation, just to name some of the important benefits.

There are hundreds of ticketing system vendors on the market today, and each one has some unique features. When choosing a ticketing system, the first decision you must make is to host it on-premises, also known as on-prem, or host it through a cloud option from the vendor. Many ticketing systems can be purchased as a software-as-a-service (SaaS) cloud model. You are charged a monthly fee based on the number of tickets, storage, or the number of agents (help desk personnel), whereas on-prem systems are licensed either per agent or per supported users, and storage is not a concern. There are a number of considerations when choosing between on-premises and cloud-based. The SaaS option is a great option for a number of reasons, mainly for availability. However, if all of your employees are local to your site, then on-premises might be the best solution. If you are conscientious about keeping the system up to date for security patches, having an on-premises system is another box to patch (so to speak). But if you elect SaaS hosting, the vendor is responsible for patching and keeping the system secure.

Besides the initial hosting options, there are several features that are somewhat standard with any ticketing system. These features include escalation, automated routing, knowledge base, email to open a ticket, ticket management, and real-time reporting, and those are just a few of the features. We will cover many of them in this section.

Entering Tickets

Once you have a ticketing system in place, it's just a matter of having tickets entered into the system. Your users should have no problem finding network issues to create tickets for. There are three main entry methods of generating a ticket: email, portal, and manually.

Email is probably the easiest way to enter a ticket. It is customary when setting up a ticketing system to dedicate an email address for entering tickets, such as `helpdesk@wiley.com` or `support@wiley.com`. When a user has a problem, they simply have to email the ticketing system, and a ticket is automatically generated. Their email address will become the identity for the ticket related to the problem, and IT can converse with them via email. This method of ticket entry is commonly used for external customers, such as product support.

Every ticketing system will have a portal for the users and help desk personnel to log into. This of course means that the users and the help desk personnel need an existing account, or users can sign up for an account when they enter the ticket. These types of ticketing systems are often established for internal issues inside the organization. Once signed into the portal, a user has several fields to fill out in order to submit the ticket. This type of setup is nice for users if they have multiple tickets and want to track them simultaneously.

The manual entry of tickets is a catch-all entry method. A help desk support person is responsible for entering the ticket information manually to create a ticket. This is a common practice when someone calls into the help desk. The entry of the ticket serves two main purposes: The first is that it allows for follow-up or escalation of the problem. The second purpose is that entering a ticket can identify a problem common to your network—for example, if a system has just been upgraded and calls come into the help desk because of a problem with the upgraded system. The user or even your help desk personnel might not know of the recent upgrade. However, the accumulation of tickets for similar issues will allow help desk personnel to identify a systemwide problem. The problem(s) can then be escalated to the party responsible for the problematic upgraded system. There are many other reasons for manually entering a ticket, such as scheduling of help desk staff, identifying specific skills required by help desk staff, and the overall volume of work. These are just a few; each organization has its own reasons.

These are the most common ticket entry methods for ticketing systems. There are many other methods, such as application-triggered entry and interactive voice response (IVR) automation.

Ticket Details

The devil is in the details, as the saying goes, and that is also true when entering a ticket into the ticketing system. There are several elements that need to be entered correctly for the successful resolution of an issue. If a user is self-entering the ticket, then you will have no control over this initial process.

When a user enters a ticket using the portal, it is common for them to be very brief with their issue description. This is usually because the description window is one field of many and the user may feel they have to be brief. Issues that are submitted via email might have more elaborate explanations, because users may feel more comfortable communicating in email. However, other fields may not be filled out, such as severity, contact information, or issue category. The ticketing system generally cannot decipher these elements from a simple email.

When acting as a help desk support person answering the help desk line, we have the most control over manually entering a ticket into the ticketing system. We also directly

interface with the user. So, there should be a standardized process for collecting the necessary information, as well as some best practices to be followed. Always keep in mind that your interactions with the user make an impression on the entire department. You should have good interpersonal skills, display empathy, and above all have patience. All of these skills will help both you and the user reach a resolution more quickly.

Another important skill that you should display is actively listening to the user. This skill is often related to the skill that detectives have when interviewing a person. When a user calls into the help desk, the critical information to collect is who they are, how to contact them, and a description of the issue. In addition, you should ask the urgency of the issue.

If you are able to enter information into the ticket entry form as the person is talking to you, do so. If you are unable to type and listen, or you must allow for long awkward pauses as you type, you should use the trusted method of pad and paper. You can always enter the information after the person is off the phone and moves on to their next task, awaiting a response for the issue. Or, if you solved the issue, then you can enter the ticket with the resolution after they are off the phone to act as a follow-up. Always exercise speed and accuracy in obtaining the information, especially if you are not able to help the person and need to escalate the ticket to someone else.

The following is a list of information you should obtain on the initial call with the user as you exercise all the best practices:

User Information The first information you want to obtain is the user's name. This is important for your communication with the user. Always use their name throughout the interview about the issue. In addition, you should obtain the best method to contact them, such as their phone number, cell phone number, email address, or physical location. This information is important if you must escalate the issue to another technician.

Description of the Issue A good description of the issue is the most important information you must gather. During your communications with the user, ask them to describe the exact issue they are having. If the user offers multiple issues, ask to review one at a time. Also, ask the user how you can reproduce the issue, then review the steps to reproduce the issue with the user. After you learn about the issue and the steps to reproduce it, repeat the steps to the user in a calm way to verify that you have clear and concise information.

Device Information Although the device being used by the user might not be apparent to you as being part of the issue, if the ticket needs to be escalated it may help in solving the issue more quickly. Information to be obtained is the type of device, operating system version, hardware service tag, and method of network access. Another technician might find that a model or version of operating system is a key to the issue.

Severity of the Issue You may pick up the urgency in the user's voice, and that can serve as an indicator of the severity of the issue. You must communicate this in entering the ticket—often there is a drop-down menu for low, mid, or high severity. Always keep in mind that not everything can be of the utmost urgency; although we would like to make everything an emergency, we often don't have the staff to do so. Communicate with the user about the severity of the issue and agree upon a severity level. Always use patience, empathy, and professionalism in your communications.

Category of the Issue The category of an issue is a critical element of the ticket. In some cases, the category may automatically route a ticket to a group of people in your support staff. For example, if a user calls in and has an issue with the organization's database and reporting, the ticket may need to be routed to the database administrators for the organization. Likewise, if there is an issue with a phone, the ticket may be routed to the phone administrators for the organization. In addition to automatic routing to a group of administrators, categories are useful for reporting. They help identify categories of issues that see the most tickets, which can be an indication you need more staff or need to adopt a quality control process.

Levels of Support

Every IT department has a structure, and they vary from organization to organization. On a very high level there are typically two main groups of support personnel: network administrators and application/database administrators. However, your organization might have security administrators, application developers, storage administrators, virtualization administrators—and these just scratch the surface.

In each group of support personnel, there are varied levels of experience, support, and responsibilities. The simplest structure is front-office personnel, who interface with the users, and back-office personnel, who interface with the front-office personnel and make systemwide changes. Depending on the size of your organization, you may also have an intermediate level of personnel who interface with the backend and frontend personnel. This intermediate level serves as a buffer to keep engineers separated from the day-to-day issues. These levels are often numbered from basic knowledge to expert knowledge, level 1 through level 3.

Regardless of your structure, these various levels are considered escalation points. As an IT technician or even an IT administrator, you are not always expected to have the answer, but you are expected to be able to get the answer. When you don't have the answer and must ask someone more knowledgeable, this process is called escalation. With a ticketing system, if all the information is properly obtained, you can escalate the ticket to a higher level of technician or administrator. The next level up the support chain should be able to read through the ticket notes and work on a resolution. When a ticket is escalated to another person, it is common for that person to own the ticket (issue) and be responsible for communicating with the user.

You may also find that you are the most knowledgeable person about a particular system in your organization, yet you still do not have an answer. This is fine and it happens all the time, but always keep in mind you that are expected to be able to find the answer. This requires an escalation of the issue to a third party, potentially outside your organization. It is also the reason support contracts should be kept current so that you have an escalation point outside of your organization. When an escalation is made outside of your organization, the point of contact (POC) inside your organization will be the owner of the ticket.

Clear Communication

The requirement and benefits of clear communications cannot be overstated enough—not just verbal communications, but also written communications between technicians and users, as well as technicians and their escalation points. Clear, concise written communications can also break down verbal communications issues between regions of the world. There are four stages to any issue where clear, concise written communications are required:

Issue Description Not only should the issue description be clear and concise to you and the user, but it should communicate the issue to someone not part of the initial communications. Don't use jargon, acronyms, or references to anything that is not clearly defined in the issue description. Always keep in mind that the issue may have to be escalated.

Work-Arounds When an issue is identified and a solution requires time, a work-around should be communicated to the user. A work-around serves as a temporary way of allowing the user to continue their work by working around the issue. Always communicate that it is a temporary solution and that you will work on a permanent solution.

Progress Notes The most important part of the ticketing process is the progress of solving the issue. You should adopt the rule of daily communications with the user and communicate where you are in the process of obtaining a resolution. Be as transparent with the user as possible—always keep in mind that the user might ask for a supervisor to review the progress. If the user stops communicating with you on the progress of a resolution, give them a drop-dead date when the ticket will be suspended or closed.

Issue Resolution The last stage to any issue is the resolution, and you should strive to have a resolution to every issue. The resolution or steps to the resolution should be clear and concise written communications. This also serves as a reference to other technicians with the same issue.

Follow-Up

Although the issue resolution should be where the communication with the user ceases, it is important to follow up with the user to make sure the issue is resolved. This important step should be done right before the ticket is closed. You may learn that the final resolution presented does not work for the user and that they continue to use the work-around that you presented as a temporary fix. This often happens when you must escalate the issue and communications are not clear and concise between the next technician and the user.

The ticket management process should include the steps of entry, resolution, solution, follow-up, and knowledge base. Each ticket that is completed strengthens your technical support for future issues. It also strengthens the faith the users have that the IT department can solve their issues. When you follow up with the user, you verify that all of your efforts and your escalation point's efforts are justified by the solution.

Many ticket systems allow for the entry of a knowledge base article for future users and technicians to self-service their issues. Writing the draft of a knowledge base article should happen after you have identified a successful follow-up with the user.

Asset Management

Asset management is an important part of the IT department's responsibilities, because the IT assets are considered *fixed tangible assets*. Some other examples of fixed tangible assets are land, furniture, and office equipment. When equipment is initially purchased, the accounting department records it as an asset on the company's general ledger, because it adds to the value of the company. Over time, however, the asset will lose its initial value. The accounting department will depreciate the value of the asset based on its perceived lifespan.

The management of these assets benefits the organization in defining the organization's worth. The management of assets also helps the IT department in forecasting upgrades and future expenditures for growth. In this section we will cover the various elements of asset management as it applies to the IT department.

Asset Management Systems

There are a number of ways to manage assets for the organization. Choosing a way to manage assets depends on what needs to be done with the information. Asset management at an organizationwide level is often a module of an accounting package used by the company. This software allows an asset (equipment) to be tracked by associating a number on the asset tag with the condition, business unit, and perceived value of the equipment. Examples of this equipment are desks, land, and even computer equipment. These types of databases work well for reporting on the value of equipment that the organization owns to calculate a net worth for an organization, but they do very little in helping an IT department plan upgrades.

Laptops, desktops, and other devices have variables such as storage, RAM, operating system versions, and other unique variables to the hardware and software of the device. Asset management systems are databases that collect data from the operating system through the use of an agent. This type of asset management is more detailed than a purchasing record from the accounting department. Once the information is collected, reports can be drawn when upgrades are required. For example, a report you may compile in the asset management system might be all operating systems that match Windows 10 and that have less than 4 GB of RAM and hard drives smaller than 100 GB. You then have a report of what needs to be upgraded in terms of hardware for an upgrade of the operating system to Windows 11. When using an asset management for an organization that spans a large geographic area, this is invaluable information that otherwise would have taken days to collect.

Software and Licensing

Asset management systems don't stop at hardware; software packages and their accompanying licensing are considered assets as well. Many asset management systems can also collect a list of the software installed on the devices in your organization. They can also include

detailed licensing usage information so that you can gauge where licensing is being used efficiently and where it is not based on usage.

Inventory Lists

Not all asset management requires databases and asset management systems. When managing a small amount of equipment, an inventory list is more than sufficient. The list can be a simple Microsoft Excel sheet detailing the types of equipment and their associated quantities. These inventory lists work really well when trying to control consumable electronics like mice, keyboards, and monitors. Once the rotating stock of equipment becomes too large in quantity and value, it's time to look at an asset management system.

Configuration Management Database (CMDB)

The configuration of devices and systems can also be assets. The fact is that your organization didn't pay for goods—they paid you for services to configure the device. The resultant configuration is considered an asset to the organization.

A configuration management database (CMDB) is a centralized repository that stores information about the configuration of items within an IT organization. These items, known as configuration items (CIs), can include hardware, software, systems, facilities, and even personnel. This centralized information can also be used for making business decisions and running efficient IT operations. Organizations can better manage changes, incidents, and compliance requirements when they know the configuration items.

The use of a CMDB enhances visibility and control over the IT environment. The CMDB helps with better incident and change management. When an incident occurs, the CMDB can be used to quickly identify the affected CIs and their relationships. The access to this information facilitates faster root cause analysis and resolution. Similarly, during change management, the CMDB provides a clear picture of the potential impact of proposed changes, helping to minimize risks and ensure successful implementation. We will cover incident and change management later in this chapter.

Asset Tags

All computer and network equipment should be tracked, from the cradle to the grave, by the IT department. When equipment enters the company, it should be labeled with an *asset tag*, as shown in Figure 10.1. The asset tag is often a permanent metallic sticker or metallic plate that is riveted to the equipment. The asset tag often has a *barcode*, which defines the numeric number that identifies the asset. This asset tag should then be entered into the asset management software, by either typing in the number or scanning the bar code.

FIGURE 10.1 An asset tag




Procurement Life Cycle

The entire life cycle for any IT system is cyclical and differs slightly depending on the assets. The typical life cycle for IT systems consists of purchasing, deploying, managing, and retiring. The exact procedures for the IT life cycle will depend on your organization and the goods or services.

The first step is the procurement of the goods or services. This process is typically standardized by your organization's business affairs department or finance department. Just like any process, the procurement process differs from organization to organization. An example of a procurement life cycle is shown in Figure 10.2; yours may differ slightly. Regardless, the procurement life cycle will always start by identifying the need for the goods or services being requested. If it costs money for the organization, then it has to fill a need or solve an issue. This is probably the most important step in the procurement process. If it does not solve an apparent issue or need, then it might not get approved or may be sidelined for another budget cycle.

FIGURE 10.2 An example of a procurement life cycle



After the need is identified, it's time to obtain quotes for the goods or services. You should have three comparative quotes for the goods or services being requested. However, depending on the goods or services, achieving that may not be feasible. Examples where comparable quotes are unobtainable are direct purchasing from the vendor, vendor-registered value-added resellers (VARs), and custom goods or services.

The budget approval process will be dependent on the reported needs of the goods or services, as previously explained. The goods or services are submitted for approval for the operational budget (OpX) or the capital budget (CapX). Items being submitted for the CapX budget will depend on their value and utility. The test is always, can the item be depreciated over the expected life of the product? Examples are servers, workstations, and other equipment. Anything that cannot be depreciated, such as services, will fall into the OpX budget. It is likely that your organization has a standard characterization of goods and services and which budget is applicable. The outcome of the budget approval process will either be approved or denied, but it can also be conditionally approved based on meeting goals or other conditions.

Once the goods or services are approved, then your business affairs department will work with the vendor or reseller to negotiate money, terms and conditions, and the overall contract/scope of work (SOW). Once the purchase is completed, you will receive the goods or services. At this point, the contract or SOW is important, because it will define when the vendor is to be paid. If all the goods are not received or the services are not complete, then the vendor is not entitled to send an invoice for payment or the invoice can be held. Although this sounds like a simple part of the procurement process, it is often overlooked. You should never begin payment until the goods are received or the contract/SOW is satisfied.

During the initial phase of obtaining quotes, you should identify the cost for ongoing support, maintenance, or licensing of the goods or services. These costs should be submitted to the operational expense budget as an ongoing/recurring cost, since these costs are usually a monthly or annual cost. The vendor or VAR might also include 3–5 years of support, maintenance, or licensing in the quote so that it can be submitted to a capital budget. This will be based on your organization's processes. Typically, the vendor will include an initial warranty or license with the original purchase of the product or service.

Every product or service outlives its usefulness. This is where we identify retiring or upgrading the product or service. This retirement or upgrade will then start the procurement process all over again. This time around, identifying the needs is easier, unless the retirement does not necessitate replacing the product.

Assigned Users

When assets are acquired by the organization, they must be managed throughout their life cycle. This typically requires assigning a person to manage the group of assets, such as laptops, servers, and hotspots. This is a critical step in the management of the asset. The person who manages the assets is responsible for identifying users who are assigned to the devices in the event of termination. The responsible person is also required to forecast upgrades and perform accounting for all assets over their life cycles. Each organization has its own requirements, but these are the top requirements for asset management.

Common Documentation

Documentation is extremely important to an IT department, not to mention the entire organization. It serves many different purposes, such as educating new IT workers, recording work performed, highlighting issues, and describing normal functionality. However,

documentation is usually one of the functions that suffer the most when projects are hurried and there is a push to start the next project.

In the following sections, we will cover the most common documents that help support the IT department and day-to-day operations. Some of these documents are prepared by collecting information in a specific manner and then detailing the results; examples are incident and standard operating procedures.

Incident Documentation

An *incident* is any event that is unusual or outside of the normal processes. You may encounter many different types of incidents as a technician: network security incidents, network outage incidents, and even customer service incidents. Regardless of which type of incident transpires, an incident document should be completed so that there is a record of the event. A record of the incident allows for further review after the incident has subsided so that it is not repeated.

The incident document should be completed as soon as possible so that key details are not forgotten. This document is often used as an executive brief for key stakeholders in the company, such as C-level people—for example, the chief information officer (CIO). The incident document can also be public-facing and used to inform customers of the incident. When used in this fashion, the incident document allows the organization to communicate with transparency about a major incident they may have experienced. Volume 2 Chapter 9, “Safety and Environmental Concerns,” covered the processes and procedures for incident response in further detail. Here are common elements of a network incident document:

Date and Time The date and time of the incident is probably the most important element, because several incidents could happen on the same day. The date and time allow you to distinguish between various incidents. The date and time should be referenced using the *Universal Time Code (UTC)* so that someone anywhere in the world does not need to calculate the time offset.

Summary The incident summary is another important piece of information that should be documented. It will often be the first thing that is looked at because it describes what happened during the incident.

Root Cause The root cause is the reason the incident occurred. Every incident needs to have a root cause defined; otherwise, the incident could happen over and over again. The network team’s first responsibility is to identify the root cause.

Actions Taken The actions taken are the actions that transpired during the incident to rectify the situation. Most of the time, they are temporary actions so that business can resume.

Remediation The remediation is the ultimate fix that will repair the root cause of the incident. This could be as simple as installing a patch, or it could be as complex as redesigning a network component.

Services Impacted The services impacted section details all the network systems and business units that were affected during the incident. This section defines the severity of the incident.

Recommendations The recommendations section details the next steps to take to remediate the issue. It also explains how to avoid similar incidents in the future.

Although these are the most common elements of an incident document, the document is not limited to these elements. Each organization has different needs for the process of reviewing network incidents. A template should be created so that there is consistency in the reporting of incidents.


Standard Operating Procedures

When organizations create policies, they will outline specific processes to adhere to the policies. Throughout this discussion you may see the words *policy* and *plan*. Policies are also considered plans for the organization; once a plan is ratified, it becomes a policy. Put another way, a policy is a mature and enforced plan of action. Each process derived from the plan or policy contains a list of steps that are called *standard operating procedures (SOPs)*, as shown in Figure 10.3. All of these components are part of the documentation process for a *quality management system (QMS)*. QMSs are created to meet certain International Organization for Standardization (ISO) requirements. A common ISO certification is ISO 9001. When a company is ISO certified, it means that it adheres to strict quality standards for consistent outcomes and has a detailed operation plan. There are many information technology ISO standards that your organization can be certified with. You have probably seen these certifications in a sales manual at some point.



ISO 9001 is a globally recognized standard for Quality Management Systems (QMS), published by the International Organization for Standardization (ISO). It helps organizations ensure they meet customer and regulatory requirements while improving customer satisfaction. The latest version, ISO 9001:2015, focuses on a process-oriented approach and risk-based thinking. Adopting ISO 9001 demonstrates a commitment to quality and continuous improvement.

FIGURE 10.3 Standard operating procedures



A process in a QMS is just that—a process. A process is defined as taking an input and creating an output. Here's an example: a specific server needs to be decommissioned, which is the input to the process of decommissioning a server, and the output of the process is that the server is wiped clean of corporate information. The process in this example is the decommissioning of a server. (We are oversimplifying the process in this example.) Most processes have an input specification and an output specification.

The SOP in this example outlines how to perform the process of decommissioning the server. The SOP should clearly explain who is responsible and the standard they must achieve for the process. In the example of a decommissioned hard drive, the SOP would define the following:

- Who is responsible for removal of the server from the rack
- Who is responsible for wiping the drives in the server
- What standard should be used to wipe the drives
- Who is responsible for removing DNS entries
- Who is responsible for removing firewall rules
- Who is responsible for clearing the BIOS/UEFI
- In which order the tasks are performed

Several tasks are created from the SOP document. Each task is defined as part of the procedure to achieve the decommissioning of the server. The work instructions serve two primary purposes. The first is to detail how each task should be performed. The exact steps are listed in the work instructions, and for the previous example, they may include the following:

- How to remove the server from the rack
- How to wipe the drives in the server
- How the drives should be wiped
- How to remove the DNS entries
- How to remove the firewall rules
- How to clear the BIOS/UEFI

The second purpose of the work instructions is to provide a training process for new employees. The work instructions are a more detailed portion of the procedure, so it becomes a training item for new employees on how to perform their job.

Procedures for Installing Software Packages

If something is worth doing the first time, then it's worth documenting and chances are it will need to be done again. This is often the case when installing custom software packages. When a new software package is purchased, a technician is often assigned to assist and install the application for the users. It is that technician's responsibility to document the software installation process. This ensures that other technicians will not have to start from scratch to help the next user who needs the software package installed.

Documenting the software installation will save time for consecutive installations. It also allows the initial technician to pass the knowledge on to other teammates. This is obviously important for the rest of the IT team, and it's also important for the initial technician. It is often the case that the original technician will be assigned the consecutive installation, mainly because they are the only ones who know how to install the software. Therefore, documenting the process of installing the software serves two purposes of saving time and passing information on to other teammates.

Onboarding/Offboarding Procedures

As employees are hired in your organization, a certain amount of initial interaction with IT is required. This interaction is called the *onboarding procedure* and is often coordinated with the HR department in your organization. During the onboarding procedure a *new-user setup checklist* should be followed. Examples of the items on the checklist include showing the user how to log in the first time and changing their password. The password policy is often the first policy discussed with the user. Other policies such as bring your own device (BYOD), acceptable use policies (AUPs), and information assurance should also be discussed during the onboarding procedure. Email, file storage, and policies should be covered as well. Each organization has a different set of criteria that make up the onboarding procedures.

Eventually, employees will leave your organization. The offboarding procedure ensures that information access is terminated when the user is terminated. A *user termination checklist* should be followed during the offboarding procedure. The process will be initiated by the HR department and should be immediately performed by the IT department. This process can be automated by using the organization's employee management system. The procedure can also be manually performed if the employee management system is not automated. However, the procedure must be performed promptly, since access to the company's information systems is the responsibility of the IT department. During the offboarding procedure, email access or BYOD access is removed through the use of the mobile device management (MDM) software; the user account is disabled; and IT should make sure the user is not connected to the IT systems remotely. The offboarding procedure may also specify that a supervisor must assume ownership of the terminated employee's voicemail, email, and files.

SLA Agreements

A *service-level agreement (SLA)* is an expected level of service that a service provider will adhere to for uptime. The SLA is detailed in the service contract as a percentage of uptime per year; it is often called "the nines." When a service provider exceeds the percentage of uptime defined in the SLA, the provider is in breach of the contract. The contract will often identify what is considered in the SLA requirement as well as the terms if they breach the SLA requirement.

You should be careful with SLAs because being up 99 percent of the time means the service provider can be down 3 days and 15 hours over the course of a year before breaching the contract. A service provider with 5 nines (99.999 percent) can be down 5 minutes and 15

seconds over the course of a year before breaching the contract. Table 10.1 details the different nines of an SLA for uptime.

TABLE 10.1 Uptime for nines of an SLA

SLA %	Downtime per year	Downtime per month	Downtime per day
99% (two nines)	3.65 days	7.2 hours	14.4 minutes
99.9% (three nines)	8.76 hours	43.8 minutes	1.44 minutes
99.99% (four nines)	52.56 minutes	4.38 minutes	8.64 seconds
99.999% (five nines)	5.26 minutes	25.9 seconds	864.3 milliseconds
99.9999% (six nines)	31.5 seconds	2.59 seconds	86.4 milliseconds

Internal

Internal service level agreements (SLAs) are used to ensure that teams within an organization maintain a specific performance goal. These agreements outline specific performance metrics, such as response times, resolution times, and quality standards, which are expected to be met by the team. By doing so, they create a clear understanding of what is required and help to manage expectations.

One of the key benefits of internal SLAs is that they promote accountability. Each team knows exactly what is expected of them and can be held responsible for meeting those expectations. This accountability helps to drive performance improvements and ensures that any issues are addressed promptly.

Knowledge Base/Articles

If time is spent on an issue, it's worth documenting so that the same amount of time is not required by someone else or yourself in the future. A *knowledge base* is a collection of issues with solutions that both your internal customers (IT staff) and external customers (end users) can use to solve common issues. The Microsoft Knowledge Base is a great example of a knowledge base. It contains more than 200,000 public articles and just as many that are private and accessible only by partners. Most helpdesk software allows for the creation of a knowledge base article from the resolution of an issue. A knowledge base can become very large, so many knowledge bases allow keyword searches.

Fortunately, you don't need fancy helpdesk software to create a knowledge base. You can simply have a collection of articles that are accessible to either your colleagues or your end users. A knowledge base article should be clear and easy to understand. Make

sure that you define any terms or jargon used within the article so that you don't lose the audience it is intended for. Here are the common elements of a typical knowledge base article:

Title The title should represent the key symptom of the issue and be action-based. For example, if the issue is that the Num Lock key stays on after the system boots up, the title could be "Num Lock Enabled After Boot."

Introduction The introduction should explain, in plain language, what the rest of the article will cover. It should include the thesis of the issue as well as the systems affected by the issue.

Symptoms The symptoms section explains how to identify the symptoms of the issue. It assists the user in validating that the issue covered in the introduction is affecting their system. This section can include steps to reproduce the issue. For example, after a fresh bootup and login, your Num Lock key stays lit up.

Solution The solution section contains the steps to fix the issue outlined in the introduction. The solution should not be restricted to just one solution; many solutions can be detailed, as well as work-arounds, if the primary solutions cannot be implemented.

More Information The more information section contains a list of resources that were used to formulate the solution. This section is important because the symptoms might be similar to those of another issue or an unidentified issue. The resources listed will allow the technician to research the issue deeper.

Date The date is important because it allows the reader to see if the issue is an old issue or a newer one, to gauge how relevant the solution is to the current issue.

Author The author section allows for the technician reading the knowledge base to contact the original technician who created the article.

Although these are the most common elements for a knowledge base article, you are not constricted to only these elements. Whichever format you end up using, it is important to be consistent. A template should be drafted so that there is consistency. Consistency allows either the technician or the end user to expect the elements when they review the article.



Exam Objective 4.2, "Given a scenario, apply change management procedures," requires you to know the change management process and how it is applied in an organization to reduce risk of changes. In addition, you should know the various components of the documented business processes, such as the rollback plan, backup plan, sandbox testing, and responsible staff members.

Change Management Best Practices

When you implement a new system or change an existing system, you affect a lot of people. You also affect business processes and other business units with these changes. Don't underestimate the power of the documentation you produce as a technician. It can and often will be used by change management groups to review the impact of your proposed changes.

Change management is a process often found in large corporations, publicly held corporations, and industries such as financial services that have regulatory requirements. However, change management is not exclusive to these types of organizations. The main purpose of change management is to standardize the methods and procedures used to handle changes in the company. These changes can be soft changes of personnel or processes, or hard changes of network services and systems.

When changes are proposed to a process or set of processes, a document is drafted called the *change management plan document*. This document is used throughout the change management process to evaluate the impact to the business continuity of the organization. In the following section, we will discuss the elements of a change management plan document.

Documented Business Processes

The documented business process is incorporated into the change management plan document. It provides an overview of the business process that the changes are expected to affect. It allows everyone involved in the process both directly and indirectly to understand the entire process.

The documentation specifically defines who interacts with, how they interact with, why they interact with, and when they interact with the process. For example, if your company created widgets, your documentation might detail the process of manufacturing the widget. The document would describe the following:

- How the raw components enter the production line
- Who delivers the raw materials
- When they deliver the raw materials
- Who assembles the raw materials to create a finished product
- Who inspects the finished product
- When and how the product is shipped

Rollback Plan (Backout Plan)

The *rollback plan*, also called the backout plan, describes the steps to roll back from a failed primary plan. If it were determined that the primary plan could not be completed, you would either implement an alternate (secondary) plan or a rollback plan, depending on the changes proposed in the change management plan document. Like the primary and alternate

plans, the rollback plan should contain the steps to be taken in the event the rollback plan must be executed. The rollback plan should also document any changes to configuration so that it can be reverted back. Most of the rollback plan will consist of the original configuration, with any additional steps to revert it back.

Backup Plan

The *backup plan* is an essential component of the change management process, safeguarding against potential failures that may arise during implementation. This plan involves creating comprehensive backups of all critical systems and data before any changes are made. The backups should be performed systematically, ensuring that each element of the infrastructure and data is covered. Verification of the backup integrity is also a crucial part of the backup plan. This involves regular testing to confirm that the backups are complete and can be successfully restored.

In the event of an unforeseen issue or failure, this plan ensures that the organization can quickly restore its systems to their previous operational state with minimal disruption. The detailed steps for restoring systems from backups should include configurations, revalidation processes, and checks to ensure operational integrity after restoration. The backup plan should be created in conjunction with the rollback plan.

Sandbox Testing

Sandbox testing is extremely useful when you want to test a change before placing it into production. A sandbox can be constructed to match your environment; you can then implement the change and fine-tune your primary plan. The use of a sandbox testing environment allows you to hone your process for the proposed change while observing any potential issues.

The introduction of virtual machines makes it very easy to set up a sandbox for testing. You can clone production servers into an isolated network and then create snapshots on the server in the sandbox and test over and over again, until all the bugs are worked out of the primary plan.

Responsible Staff Members

Every process in the organization must have a person who is assigned to be the *responsible staff member*. This person oversees the process and can answer questions about the process. If there are any changes to the process or changes that can affect the process, this person acts as the main point of contact. They can then facilitate any changes to the process.

As an example, you may assign a person to be the responsible party for the electronics decommissioning process. Any questions about disposal of electronics should be directed to this person. If your organization is choosing a new e-waste company, it will affect the decommissioning process. Therefore, this person should be included in the decision as a stakeholder. Any changes can then be adjusted or integrated into the decommissioning process, and this person can facilitate the changes.

Request Forms

The change management process often begins with a *request form* that details the proposed change. The exact elements in the request form will differ slightly depending on your organization's requirements. The following lists the most common elements found on the change management request form. Some of the information found on the request form is preliminary; the information will be expanded upon as the request form transitions into the change control document.

Item to Be Changed This is the item that is being requested for a change. Examples are software, hardware, firmware, configuration, or documentation.

Reason The reason the item is being submitted for a change. Examples are legal, marketing, performance, software bug, or process problem.

Priority The urgency, or priority, of the change is documented in this section of the form, although this should also be conveyed in the reason for the change. The priority is often a separate field on the change control form. Examples are emergency, urgent, routine, or a specific date.

Change Description/Plan The description or plan for the change is documented in this section. For changes in configuration, you would detail the changes to the configuration and why each part of the configuration is being changed. Firmware changes would list the version being upgraded from and the version being upgraded to.

Change Rollback Plan The rollback plan describes the steps to roll back from a failed primary plan. If it was determined that the primary plan could not be completed, you would implement either an alternate plan or a rollback plan depending on the changes proposed.

Technical Evaluation In this section of the form, you document why the primary plan will succeed. The changes should be tested in a lab environment closest to the production environment and documented in this section. When you're creating the technical evaluation, outline specific objective goals along with the metrics with which they can be measured.

Duration of Changes Here, you document the estimated duration of the change process. Any service outages will be documented in this section.

Purpose of Change

The *purpose of change* is the reason the change management process starts. Either your business unit requires a change that will affect others, or another business unit requires a change that can affect your business unit indirectly. A change to any part of the process, such as the intake of raw materials, could affect the end result. Change is an essential component of a business, and it should be expected. If your company only created widgets and never evolved, you would eventually be out of business.

Unfortunately, not all changes support the company's product line directly. Some changes are imposed on the company, because IT systems are constantly changing. As a technician,

you are constantly upgrading and patching systems and equipment. These upgrades and patches are considered changes that can affect the entire business process.

This section of the change management plan document should explain why the change is necessary. It needs to include any vendor documentation explaining the change to the product. For example, if the proposed change were to install a Microsoft Windows security patch, the purpose of the change would be the security of the Windows operating system. The vendor documentation in this example would be the knowledge base article that normally accompanies Windows security patches. Other examples of purposes of change might be legal, marketing, performance, capacity, a software bug, or a process problem that requires a change.

Scope of Change

The *scope of change* details how many systems the proposed change will affect. The scope could involve only one system, or it could be all the systems in an entire enterprise. The scope of change is not limited to the number of systems that will be changed. The scope can also describe how many people a proposed change will affect. For example, if you propose to change lines of code in an ordering system, the change could affect your salespeople, customers, and delivery of the products. The scope of this change could impact the business continuity directly if something goes wrong during the change.

When creating this section of the change management plan documentation, be sure to document which systems the proposed change will affect, the number of systems the proposed change will affect, the number of people the proposed change will affect, as well as whether anyone will be directly or indirectly affected. In addition, you should include the proposed date and time of the change and how long the change will take. Keep in mind that this section allows the change management team to evaluate how big the proposed change is. The scope should answer the following questions:

- Who will the change affect?
 - IT resources
 - Other departments
 - Business units
 - Partners
 - Customers
- What will the change affect?
 - IT systems
 - Sites/locations
 - Processes/procedures
 - Availability of systems
 - Outage duration

- What will be impacted?
 - Other proposed changes
 - System performance
 - System capacity
 - Other resources (personnel, security, etc.)
- When is the proposed change?
 - Date and time of the change
 - Duration of the change

Change Type and Schedule

In addition to the purpose of the change, and the preceding questions about the change, the driving factor of the change is going to be the change type. The change type will ultimately affect the schedule of the change. The change type is the driver behind the change. The common change types are as follows:

Standard Change In the realm of change management, a standard change is a pre-approved, low-risk change that follows a predefined process and is frequently performed. These changes typically have minimal impact on the organization and require no additional approvals.

Normal Change A normal change is any change that does not fall under the standard or emergency categories. It requires a thorough review and approval process, including a risk assessment and a detailed implementation plan. This type ensures that all potential implications are considered before proceeding.

Emergency Change An emergency change is a high-priority change that must be implemented immediately to address critical issues that could cause significant disruption or harm to the organization. Due to its urgent nature, the usual review and approval processes are expedited, with the focus being on resolving the issue as quickly as possible to restore normal operations.

Each of these change types requires a tailored approach to ensure that they are managed effectively and align with the organization's overall change management strategy.

Maintenance Window

The change type will obviously dictate the date and time of the change or timing of the change. When deciding on the timing of the change, the maintenance window should be considered. The maintenance window is a scheduled period during which regular maintenance tasks, updates, and system enhancements are performed on an organization's IT infrastructure. This predefined time frame is carefully selected to minimize disruption to normal business operations, typically occurring during off-peak hours when system usage is at its lowest.

Change Freeze

A change freeze is a period during which no changes or very limited changes are allowed to be made to the system. This is typically enforced during critical business periods, major projects, or before and during the implementation of significant IT changes. A change freeze is a deliberate pause in the implementation of new changes within an IT environment. During the change freeze, only emergency changes or pre-approved exceptions are allowed, and these are closely monitored and controlled.

During a change freeze, the primary objective is to maintain the stability of the IT environment. With no new changes being introduced, the risk of unforeseen issues and disruptions decreases significantly. This is especially critical during high-stakes periods such as end-of-quarter financial reporting, major product launches, or other times when system availability and reliability are paramount.

Peer Review

Typically, we will not solely understand the impact of a change to other systems. This is why a peer review should be performed with peers in other departments. In some circumstances, the change might appear to be successful but inadvertently cause problems elsewhere in the business process. An example is a change made to a shipping system, which could affect your data analytics or accounting team's business process. This is where changes should be identified for risk and introduced into the risk analysis process.

Risk Analysis

Whenever a change is made to a system or equipment, there is the potential for the system or equipment to fail. The change could even cause another system or piece of equipment to fail. In some circumstances, the change might be successful but inadvertently cause problems elsewhere in the business process. For example, if a change to an ordering system causes confusion in the ordering process, sales might be inadvertently lost.

Risk analysis is the process of analyzing the proposed changes for the possibility of failure or undesirable consequences. Although you will include this section in the initial change management plan document, your sole risk analysis will be narrow in perspective, because you will focus on the process from the IT aspect. A change advisory board will perform a much larger risk analysis. This team will have a much larger perspective, since they come from various business units in the organization. From this analysis, a proper risk level to the organization can be determined. The risk level will dictate how much time is spent on the possibility of failure or undesirable consequences from the change.

Risk Levels

Risk levels are used to communicate our perceived risk for failure or a change. For most purposes the risk levels can be kept simple, such as a three-tier system of: low, moderate, or high. You can build risk level structures with more than three tiers of risk, such as the popular four-tier model of low, moderate, high, and extreme, or the six-tier model of very-low, low, medium, high, very-high, and extreme. The number of risk levels your organization

will use to convey risk will be based on several factors, such as the typical audience and the amount of quantitative data used to determine the perceived risk.

If you are analyzing risk for a development department, then you may want to use a granular system that has many tiers. Developers are accustomed to details, and they typically prioritize their work on a granular level. However, if you are presenting to a C-suite level of individuals on a daily basis, then a more simplistic system is in order, since executives drive their decisions based on the brevity of information and not always the finite details.

The three-tier model serves many organizations just fine when it comes to risk analysis. Here are the risk levels when we classify a change:

- Low Risk
 - The system is easily recoverable and reproducible.
 - The system provides an informational (non-critical) service.
- Moderate Risk
 - The system is integral with other systems.
 - The system provides a normal or important service.
- High Risk
 - The system is highly integrated with other systems.
 - The system provides a critical organizationwide service.

Implementation

The implementation section of the change management plan explains how the proposed change will be executed. Steps should be detailed on the changes and the order of the changes. If changes were to be made in configuration files, switches, or routers, you would document the changes to the configuration and why each part of the configuration is being changed. Firmware changes would list the version being upgraded from and the version being upgraded to. The idea is to provide as much detail as possible about the documented changes to be made to the systems or equipment.

When a change is implemented or planned, there is always the potential for problems, or you may identify a consideration in the execution of the plan. The plan for change section should detail those considerations. It's common for a primary plan to be drafted as well as an alternate plan in the event the primary plan cannot be executed. For example, if the primary plan is to move a server from one rack to another so that it can be connected to a particular switch, the alternate plan could be to leave it in the rack and use longer cables. Be sure to have multiple plans; once the change is approved, the plan(s) outlined in this document must be executed closely.

You should also document why the primary plan will succeed. The changes should be tested in a lab environment closest to the production environment (if possible) and documented in this section as well. When creating the plan, you should outline specific, objective goals, along with the metrics with which they can be measured. For example, if you are planning to make a change because there is a high error rate on an interface, then the metric measure to be

compared would be the error rate on the interface. You would document what you expect the error rate to be after the change is made so that you can measure the success of the change.

Change Board

The change board, also known as the *change advisory board*, is the body of users who will ultimately evaluate and then approve or deny the change you propose. This group of people often meets weekly to discuss the changes detailed in the change management plan documents. The goal of the change advisory board is to evaluate the proposed changes in order to reduce the impact on day-to-day operations of the organization.

It is common practice for the meetings of the change advisory board to be held via a conference call at a set time every week. This allows key stakeholders in an organization to be available regardless of where they are in the world. Because it's at a set time every week, there are no excuses for not being available during the change control meetings.

It is likely that if you are the technician proposing the change, you will be on the call for questions or clarification. The key to getting a change approved is to know your audience and communicate clearly in the change control plan document. Remember, the change advisory board is often composed of various stakeholders from the entire organization, not just IT. You should not assume that the change you are proposing is as clear to them as it is to you. Some change advisory boards are made up strictly of IT stakeholders, so you must understand who will review the proposed changes and choose your wording appropriately.

The change management document must be approved by the majority of change advisory board members or by specific board members. The approval of the proposed change should be documented in the change control policy for the organization. Only approved changes can be executed. If other changes need to be made outside of the original submission, additional approvals must be acquired.



Each organization is different, and each change management documentation process will be slightly different. However, the sections described here are the most common elements found on these documents.

End-User Acceptance

Although the CompTIA A+ exam does not focus on application development testing and approval, *user acceptance* is an objective on the exam as it pertains to the change management process. It should be noted that user acceptance is not solely used for application development; it is also used when there is a significant update to an interface or a process, such as a service pack or upgrade to an operating system.

When a change is to be made in which the user's interaction will be impacted, it is common practice to beta-test the change. This is also known as user testing or just plain application testing. You can achieve user acceptance two different ways:

In-Person Testing The person responsible for testing dedicates time and performs testing in person with the developer.

Self-Paced Testing This has its advantages, since the user will be relaxed in their own environment during testing.

Regardless of which method of testing you choose, a strict time frame must be communicated to the user testing the change.

Once user acceptance is obtained, it should be documented in the user acceptance section of the change management documentation. The methods of testing, the users and groups involved in testing, and the time invested in testing should be included in this section as well. Remember that the goal is the approval and successful implementation of the changes, so it is important that you are convincing and, more importantly, convinced that the change will succeed without repercussions.



Exam Objective 4.3, “Given a scenario, implement workstation backup and recovery methods,” requires you to know how to implement backup and recovery methods, backup testing, and backup rotational schemes. This will ensure you can recover from information that is deleted or when you fall prey to a disaster.

Backup and Recovery

As a technician you’re responsible for recovering from uncontrolled disasters. Luckily, you can prevent disasters by taking proper precautions, as we will discuss in this section. You’ll find that you’re prepared when disaster strikes and can restore business continuity that much more quickly. This section discusses data loss from the following types of disasters: data loss, human error, and data loss as the result of threat actors.

When we think of data backups, we usually relate them to disasters. However, data backups are not just used to restore from disaster; we often use data backups when a user inadvertently deletes files they shouldn’t have deleted. Data backups are also used when users overwrite files or just plain forget where they put them in the first place. Regardless of how the data was lost, the underlying reason we create data backups is to recover from data loss.

Because you can’t choose the disaster or situation that causes the loss of data, you should adopt a layered strategy, starting with the user and expanding outward to the infrastructure. The following sections cover several types of strategies that can protect you from data loss.

File-Level Backups


Most of the time, your users will need to restore a single file or perhaps a few files, but definitely not the entire server or server farm. Therefore, you should make sure that one of the layers of protection allows for the restoration of individual files. You can implement this type of strategy in several different ways. Depending on your resources, you should use them all.

Volume Shadow Copy

Volume Shadow Copy, also known as the *Volume Snapshot Service (VSS)*, has been an integral part of the Windows Server operating system since the release of Windows 2000.

Volume Shadow Copy can be enabled on a volume-by-volume basis. Once it's turned on, all the shares on the volume are protected. You can access Volume Shadow Copy by right-clicking a volume and selecting Properties. You can then configure it by using the Shadow Copies tab, as shown in Figure 10.4.

FIGURE 10.4 The Shadow Copies tab




Volume Shadow Copy has one amazing advantage: it empowers the user to restore their own files. All the user needs to do is right-click the file or empty space in the shared folder, select Properties, and then in the Properties window, select the Previous Versions tab. This will open a list of snapshots, as shown in Figure 10.5. The user can then double-click the snapshots to open them as if they were currently on the filesystem. This allows the user to evaluate what they are looking for. Once they find what they are looking for, they can either click the Restore button or drag the files over to the current folder.

One limitation to Volume Shadow Copies is the number of snapshots that can be active. Only 64 snapshots can be active at one time. The oldest snapshot is deleted when a new snapshot is created to maintain a running total of 64 snapshots. By default, Volume Shadow Copy is not enabled. When it is enabled, the default schedule creates a snapshot twice a day, at 7 a.m. and 12 p.m. It's advisable to set a schedule that creates a snapshot every hour

during normal business hours. This will give the user the last 64 hours of work, which could be well over a week and a half, if you were open 9–5.


FIGURE 10.5 The Previous Versions tab



File-Based Backups

File-based backups are a common type of backup in organizations today and have been since the introduction of backup software. The Windows Server operating system includes a backup program capable of protecting the local server, as shown in Figure 10.6. It is somewhat limited, because it only supports a file-based destination and does not offer options for data tapes. It also only allows for the management of the local server. However, the product is free and is included with the Server operating system, so there is no reason not to have some type of backup.

Advanced backup software, such as Veeam Backup & Replication and Veritas Backup Exec, allows for the centralized management of all backups. Multiple backup jobs can be created for various groups of servers and can be directed to various destinations. For example, the accounting servers might back up to a tape library unit, whereas the sales servers back up to a disk storage unit. We'll discuss media types later in this chapter, but the key takeaway is that multiple jobs can be created and executed at the same time.

FIGURE 10.6 Windows Server Backup

Advanced backup software often requires a licensed agent to be installed on each server. Depending on the type of agent purchased, the agent might just allow for a simple backup of files, or it might allow for open files to be backed up while they are in use. Some agents even allow for the snapshot of all files so that a point-in-time image can be made of the filesystem. The backup is then created from the snapshot. This type of backup is common in financial institutions, where an end-of-day cutoff needs to be created.

Advanced backup software normally performs a pull of files from the selected source server and directs the information to the selected media. This is called the pull backup method, and it is probably the most common type of backup you will encounter. However, there are also push backup methods, in which the backup software directs the selected source server to push the files to the destination media using the backup server. This reduces the utilization on the backup server and speeds up the backup process, also known as the backup window.

Image-Based Backups

Image-based backups allow for a complete server to be backed up. This type of backup is also called a *bare-metal backup*. It's called a bare-metal backup because if the server hardware were to fail, you would restore the backup to a new server (bare-metal) and restore it completely. The inherent problem with these types of restorations is that they require administrator intervention. However, the technology is impressive and spares you from reinstalling the server from scratch.

Virtualized environments are where image-based backups really add value. Virtualization is changing the landscape of IT, and the area of backups is no different. When a server is virtualized, the guest virtual machine consists of configuration files, a virtual filesystem file, and other supporting files. When access is given to the underlying filesystem where the files can

be directly accessed, they can be backed up. This allows for an image to be created for the current state of an operating system—files and all.

Most enterprise backup software supports image-based backups for an additional license fee. It normally requires an agent to be installed on the host operating system, such as Microsoft Hyper-V. In VMware environments, a VMware Consolidated Backup (VCB) proxy is required. This application proxy allows the backup software to create snapshots for the guest virtual machines and assists in backing up the virtual machine files.

Critical Applications

So far, we've discussed how to use file server backups to protect an organization. However, an organization does not rely solely on file servers; there are many other types of servers in an organization. Examples include Microsoft SQL, for databases, and Microsoft Exchange, for email. In addition, there are several other types of applications that might be custom to an organization.

Just like file servers, Microsoft SQL and Microsoft Exchange have custom agents that are licensed. These agents allow for the data contained in the proprietary data stores to be backed up to your backup media. In addition to the backup of data, the agent starts a maintenance process at the end of a backup. This maintenance process checks the consistency of the current data store by replaying transaction logs, also called *tlogs*.

Critical applications for an organization do not have to be on-site. As organizations adopt a cloud-based approach to IT, they push critical applications out of the network and into the cloud. Providers such as Amazon Web Services (AWS) and Microsoft Azure offer not only the critical applications but also backup services that are contained in the cloud.

Backup Media

When creating a backup job, you choose what you want to back up (source) and a destination where it is to be stored. Depending on the backup software, you may have several different destinations to select from. Examples include iSCSI storage area networks (SANs), network-attached storage (NAS), tape library units (TLUs), or even cloud-based storage, such as Amazon S3. These are just some examples; there are many different media options for storing backups. Each backup media option uses a specific media type, and each media type has unique advantages and disadvantages. Here are the three media types commonly used for backups:

Disk-to-Tape Disk-to-tape backups have evolved quite a bit throughout the years. Today, *Linear Tape-Open (LTO)* technology has become the successor for backups. LTO can provide 6 TB of raw capacity per tape, with plans for 48 TB per tape in the near future. Tapes are portable enough to rotate off-site for safekeeping. However, they require time to record the data, resulting in lengthy overall backup time. Restore requires time to tension the tape, locate the data, and restore the data, making the RTO a lengthy process.

Disk-to-Disk Disk-to-disk backups have become a standard in data centers as well because of the proximity of the data and the short RTO. This type of media is usually based on-site and then it is used to create an off-site copy. It can record the data faster than traditional tape, thus shortening overall backup time. It also does not require tensioning and seeking for the data, like a tape requires.

The capacity of a disk, however, is much smaller than a tape because the drives remain in the backup unit. Data deduplication can provide a nominal 10:1 compression ratio, depending on the data. This means that 10 TB of data can be compressed on 1 TB of disk storage. So, a 10 TB storage unit could potentially back up 100 TB of data. Again, this depends on the types of files you are backing up. The more similar the data, the better the compression ratio.

Disk-to-Cloud Disk-to-cloud is another popular and emerging backup technology. It is often used with disk-to-disk backups to provide an off-site storage location for end-of-week backups or monthly backups. The two disadvantages of a disk-to-cloud backup are the ongoing cost and the lengthy RTO. Anytime you want to restore data, you must download the data from the cloud. Depending on your Internet connection speeds, this could require some time. The advantage is that expensive backup equipment does not need to be purchased along with the ongoing purchase of tapes.

Media Rotation

Administrators will adopt a rotation schedule for long-term archiving of data. The most popular backup rotation is *grandfather, father, son (GFS)*. The GFS rotation defines how tapes are rotated on a *first-in, first-out (FIFO)* basis. One of the daily backups will become the weekly backup on a FIFO basis. And lastly, one of the weekly backups will become the month-end backup. Policies should be created such as retaining 6 daily backups, 4 weekly backups, and 12 monthly backups. As you progress further away from the first six days, the RPO jumps to a weekly basis, then to a monthly basis. The benefit is that you can retain data over a longer period of time with the same number of tapes.

Backups are created for one of two main reasons: accidental deletion and disaster. Therefore, it makes sense that a disaster that could destroy your data center could also destroy the backup media. For this reason, media should be rotated off-site from the on-site presence of the original media.

The *3-2-1 backup rule* method is a common method for maintaining both on-site and off-site backups. The 3-2-1 method works like this: Three instances of the data should exist at all times. The original copy of the files and a backup of the files should be on-site, and the third copy of the data should be off-site in the event of tragedy at the site. Here's an example: You create a business proposal on your computer (first instance), and nightly your files are backed up (second instance). You now have two instances local to your immediate site (on-site) in the event of an accidental deletion. A second backup job then backs the file up to the cloud. This provides a third instance of the file, which is off-site.

There are a number of ways you can achieve this method of disaster recovery. For instance, you create the file, Volume Shadow Copy snapshots the drive on the hour, and a nightly backup copies the file to the cloud for off-site storage.

Backup Methods

There are several options for creating file-based backup jobs. Each backup method has advantages and disadvantages, depending on the media you are using and the amount of time in your backup window. The following are several of the backup methods you will find primarily with file-based backups:

Full A full backup is just that: a full backup of the entire data set. A full backup requires the longest backup window of all the methods, because it is the slowest. Therefore, full backups are generally performed on weekends, when you have a lengthy window. All files are backed up, regardless of the state of the archive bit. The archive bit is an attribute of each file; when a file is created or modified, the archive bit is turned on. When a full backup is performed, all the archive bits are reset on the files in the data set. A full backup is not sustainable through the week because it backs up all the files, regardless of whether they have been modified.

Incremental An incremental backup is used to speed up backups through the week when backup windows are short. An incremental backup backs up all files with the archive bit set. After the files are backed up, the archive bit is reset. Only the files that were created and modified from the last full or prior incremental backup are backed up, so backups are small. If you need to restore from an incremental backup, however, you will need to restore the full backup as well as all of the incremental backup files up to the RPO required. This type of restore creates a longer RTO because of the multiple backups that are required.

Differential A differential backup is also used to speed up backups through the week. It will back up all the files with the archive bit set as well, but it will not reset the archive bit after it has done so. A differential backup creates a gradually larger backup until a full backup is completed and the archive bits are reset again. This type of backup has a shorter RTO than incremental backups, because only the full and the last differential backups are needed to restore to the RPO required.

Copy A copy backup is used when you want to make an exact copy of the server. This backup method is identical to a full backup, with the exception that it does not reset archive bits. This backup method is often used when you want to make a copy of a server before a major upgrade. This backup method allows you to create a backup without affecting an ongoing nightly backup.

Synthetic Full A *synthetic full backup*, also known as just a synthetic backup, is a full backup without the overhead of performing a full backup on the entire data set. A synthetic backup is performed by the backup software. The software takes the latest full backups and applies the differential backups or incremental backups up to the desired date the job runs. For example, if you create a full backup once a month but want weekly full backups, a synthetic full backup can be created by taking the last full backup and applying the changes up to the current date. The synthetic full backup now serves as a full backup in the event a restore is required, thus lowering the RTO.

Data Restoration

When you have implemented the appropriate backup types for your organization, implemented the proper methods and media, and have a backup rotation, then you are prepared for when disaster strikes! If a user suffers data loss, someone deletes something inadvertently, or a threat actor ransoms data, you have the ability to restore the data. The question is: where to restore the data?

In-Place If the restore is being performed because a user inadvertently deleted a folder or folder structure, then the answer is pretty obvious, the data should be restored in-place. This may be the server share it was originally shared from. The location might also be a folder on the server that is used for application data, such as a database.

Overwrite There are certain situations where you may want to overwrite the existing data. These situations might be from ransomware, or just plain old people being people and corrupting their files. In either case, the restoration of data will need to be done over the existing data. The location could be to a share or file location.

Alternative Location When you restore data, most backup software will allow you to choose an alternative location. Using an alternative location is your best choice if a user is questioning the validity of their existing data. They might want to compare a group of files from a prior time period, against their current files. It is also a very low risk to restore files to an alternative location. However, the consideration should be the size of the files being restored and the potential for someone to start using the restored files in lieu of the current files. Therefore, communication is key with the user requesting the files to be restored. The alternative location should be labeled as a restore from a specific date.

Characteristics

When discussing the restoration of data, two characteristics dictate when you back up and how you back up. The concept of the *recovery point objective (RPO)* defines the point in time that you can restore to in the event of a disaster. The RPO is often the night before, since backup windows are often scheduled at night. The concept of the *recovery time objective (RTO)* defines how fast you can restore the data.

Backup Testing

Over the years, we've seen fellow administrators rely on their backups—up to the point when they try to restore them. It's a very different story when they fail during a critical moment. Fortunately, this only happens to you once, and then you adopt testing strategies. You should not consider data on a backup to be safe until you have proven that it has been restored successfully. There are so many things that can go wrong with a restore, the most common being media failure.

We recommend that the frequency of performing a restore of your backup is at least once a month. This will allow you to verify that you actually have data that is restorable in the event of an emergency. Many backup products restore the data to a temporary location and compares the restore to what was originally backed up on the media. This functionality allow you to schedule a test restore. When it's done testing the restore, it deletes the restored data and notifies you of any discrepancies.



Exam Objective 4.7, “Given a scenario, use proper communication techniques and professionalism,” requires you to know how to demonstrate professionalism. This is to include your proper attire, language, and attitude. You should also know how to deal with customers, handle confidentiality, and meet expectations of the customer.

Demonstrating Professionalism

As a professional technician, you need to possess a certain level of technical competence, or you'll quickly find yourself looking for work. Technical ability alone isn't enough, though; there are many people out there with skills similar to yours. One thing that can set you apart is acting like a true professional and building a solid reputation. As the noted investor Warren Buffett said, “It takes 20 years to build a reputation and 5 minutes to ruin it. If you think about that, you'll do things differently.”

You could probably break down professionalism 100 different ways. For the A+ 220-1202 exam, and for the purposes of this chapter, we're going to break it down into two critical parts: communication and behavior.

Good communication includes listening to what the user or manager or developer is telling you and making certain that you understand completely. Approximately half of all communication should be listening. That a user or customer may not fully understand the terminology or concepts doesn't mean they don't have a real problem that needs to be addressed. Therefore, you must be skilled not only at listening but also at translating.

Professional behavior encompasses politeness, guidance, punctuality, and accountability. Always treat the customer with the same respect and empathy that you would expect if the situation were reversed. Likewise, guide the customer through the problem and the explanation. Tell them what has caused the problem they are currently experiencing and the best solution for preventing it from recurring in the future.

Professional Appearance

Demonstrating professionalism begins with your appearance and attire. You should always dress for the respect and professionalism that you deserve. It's an easy element of being professional, but it is often an overlooked element.

If you dress down, then you will be judged as being less professional than your intellect or position deserves. The customer might not communicate with you the way you'd expect, because the customer will make assumptions about your intellect or position based on your appearance. The opposite can also happen if you overdress. The customer might not give you the same respect that a technician of your caliber might expect. You could be viewed as over-qualified for the customer's needs or nontechnical.

These judgments shouldn't happen from the customer's perspective, but they often do because of the impression that appearance makes. Therefore, you should always attempt to match the required attire of the environment you are working in. For example, a technician on a construction site should have the expected attire, such as safety equipment and rugged clothes. Showing up in a suit is not what the customer would expect. Conversely, if you showed up at an office environment wearing jeans and a t-shirt you'd be looked at as being unprofessional.

Luckily for IT professionals, there are two norms of appearance and attire that are expected, depending on the environment and workplace:

Formal Business formal, also known as business professional, typically requires that male employees be dressed in a suit and tie. The male attire will consist of a long-sleeve dress shirt, tie, dark or striped pants, a suit jacket, and dress shoes. Female employees should be dressed in a business suit or business skirt. Women's clothes can be lighter in color than men's clothes, and long sleeves are preferable. Women should wear heeled shoes or dress shoes.

Business Casual Business casual is often interpreted as casual and therefore many professional organizations have not embraced it. The standard definition is dress pants or khakis, collared shirt, blazer, and professional-looking shoes. For male employees a tie is not required and the color palette can be lighter for both pants and blazer. Female employees can wear a collared or noncollared blouse, dress or skirt (typically below the knee), along with heels, dress shoes, dress boots, or flats.

The dress attire of the organization will differ slightly from the preceding definitions based on season, organization type, and even the day of the week. It is popular now for organizations to have a casual Friday that differs in definition based on the norm of business formal or business casual the rest of the week. Therefore, it is always best to check with your supervisor or coworkers about what is appropriate and what is not.

Communicating with Customers

The act of diagnosis starts with the art of customer relations. Go to the customer with an attitude of trust. Believe what the customer is saying. At the same time, retain an attitude of hidden skepticism; *don't* believe that the customer has told you everything. This attitude of hidden skepticism is not the same as distrust. Just remember that what you hear isn't always the whole story; customers may inadvertently forget to provide some crucial details.



One of the best ways to become proficient in communicating with customers is to put yourself in the shoes of the novice user. None of us are experts in every field, so think of an area where you are weak—auto repair or home repair, for example—and imagine how you would want a professional in that area to communicate with you.

For example, a customer may complain that their optical drive doesn't work. What they fail to mention is that it has never worked and that they installed it. On examining the machine, you realize that they mounted it with screws that are too long, preventing the tray from ejecting properly.

Here are a few suggestions for making your communication with the customer easier:

Use proper language. End users are likely to be less computer literate than you—that's why you are there to fix the problem. Therefore, they might not know computer acronyms, slang, or jargon. In fact, the excessive use of computer terms might make them feel stupid (and make the situation uncomfortable) or make you seem aloof. It's always best to start off with basic terms, such as "Where does it plug into?" If the user answers, "Oh, it's an eSATA SSD," then you know you have the freedom to use more technical terms. Users with computer knowledge will often let you know. Users with little knowledge may be afraid to tell you that they don't understand something because they don't want to look dumb.

Have the customer reproduce the error. The most important part of this step is to have the customer show you what the problem is. The best method we've seen of doing this is to ask, "Show me what 'not working' looks like." That way, you see the conditions and methods under which the problem occurs. The problem may be a simple matter of doing an operation incorrectly or performing the operation in the wrong order. During this step, you have the opportunity to observe how the problem occurs, so pay attention.

Identify recent changes. The user can give you vital information. The most important question is, "What changed?" Problems don't usually come out of nowhere. Was a new piece of hardware or software added? Did the user drop some equipment? Was there a power outage or a storm? These are the types of questions that you can ask a user when trying to find out what is different. If nothing changed, at least outwardly, then what was going on at the time of the failure? Can the problem be reproduced? Is there a work-around? The point here is to raise as many questions as you need to ask in order to pinpoint the source of the trouble.

Use the collected information. Once the problem or problems have been clearly identified, your next step is to isolate possible causes. If the problem cannot be clearly identified, then further tests will be necessary. A common technique for hardware and software problems alike is to strip the system down to bare-bones basics. In a hardware situation, this could mean removing all interface cards except those absolutely required for the system to operate. In a software situation, this may mean disabling elements within Device Manager.

Generally, then you can gradually rebuild the system toward the point where the trouble started. When you reintroduce a component and the problem reappears, you know that component is the one causing the problem.

Customer satisfaction goes a long way toward generating repeat business. If you can *meet* the customer's expectations, you will most likely hear from them again when another problem arises. However, if you can *exceed* the customer's expectations, you can almost guarantee that they will call you the next time a problem arises.



Real World Scenario

Communication Is Key

Marriages disintegrate when couples do not communicate effectively—or so many experts proclaim. Communication is ranked as one of the most important skills needed to make a marriage work. The same can be said for business partnerships. It is important to make certain that you are listening to your customers, whether they are truly customers in the traditional sense of the word or internal users that you support. You also need to listen carefully to managers and vendors and make sure that you understand them before beginning a project.

Similarly, you need to make certain that the parties in question understand what you are saying to them. It isn't acceptable to resort to the "But I told you . . ." excuse when customers or partners aren't pleased with the results. Making certain that they understand what you are telling them is as important as making certain that you understand what they are telling you.

Customer satisfaction is important in all communications media—whether you are on-site, providing phone support, or communicating through email or other correspondence. If you are on-site, follow these rules:

- When you arrive, immediately look for the person (user, manager, administrator, and so on) who is affected by the problem. Make sure that the user knows you are there and assure that person that you will do all you can. Project a positive attitude and confidence that you can resolve their problem.
- Listen intently to what your customer is saying and avoid interrupting them. Make it obvious that you are listening and respecting what they are telling you.
- If there is a problem with understanding the client, go to whatever lengths you need to in order to remedy the situation. Look for verbal and nonverbal cues that can help you isolate the problem. Taking notes is also helpful. Make sure to say something to the user like, "I'm listening, and I am going to write down a few notes just to make sure I remember everything."

- Share the customer's sense of urgency. What may seem like a small problem to you can appear to your customer as if the whole world is collapsing around them.
- Be honest and fair with the customer and try to establish a personal rapport. Explain what the problem is, what you believe the cause is, and what can be done in the future to prevent the problem from recurring.
- Handle complaints as professionally as possible. Accept responsibility for errors that may have occurred on your part and never try to pass the blame elsewhere. Avoid arguing with a customer; it serves no purpose. Resolve the customer's anger with as little conflict as possible. Remember, the goal is to keep the customer, not to win an argument.

When you finish a job, notify the user that you have finished. Make every attempt to find the user and inform them of the resolution. If you cannot find the customer, leave a note explaining the resolution.

You should also leave a means by which the customer can contact you if they have any questions about the resolution or a related problem. In most cases, you should leave your business number and, if applicable, your cell phone number, in case the customer needs to contact you after hours.

You should also notify both your manager and the user's manager that the job has been completed.

If you are providing phone support, keep the following guidelines in mind:

- Always answer the telephone in a professional manner, announcing the name of the company and yourself.
- Using the customer's name can help build rapport. Using it in every sentence can sound condescending, but using it once in a while can make you seem more personable.
- Make a concentrated effort to ascertain the customer's technical level and communicate at that level, not above or below it.

The most important skill that you can have is the ability to listen. You have to rely on the customer to describe the problem accurately. They cannot do that if you are second-guessing or jumping to conclusions before the whole story is told. Ask broad questions to begin, and then narrow them down to help isolate the problem.

It is your job to help extract the description of the problem from the user. For example, you might ask the following questions:

- Is the printer plugged in?
- Is it online?
- Are there any lights flashing on it?

Complaints should be handled in the same manner in which they would be handled if you were on-site. Make your best effort to resolve the problem and not argue. Again, your primary goal is to keep the customer.

Close the incident only when the customer is satisfied that the solution is the correct one and the problem has gone away.

End the phone call in a courteous manner. Thanking the customer for the opportunity to serve them is often the best way.

Talking to the user is an important first step in the troubleshooting process. Your first contact with a computer that has a problem is usually through the customer, either directly or by way of a work order that contains the user's complaint. Often, the complaint is something straightforward, such as, "There's smoke coming from the back of my monitor." At other times, the problem is complex, and the customer does not mention everything that has been going wrong. Regardless of the situation, always approach it calmly and professionally, and remember that you get only one chance to make a good first impression.



Real World Scenario

Communication Is Everywhere

Almost every profession stresses the importance of good communication, and the IT industry is no different. Jamie Walters, founder and chief vision and strategy officer for Ivy Sea, Inc., and Sarah Fenson, Ivy Sea's guide to client services, wrote an article for Inc. com on steps to smooth conversations (www.inc.com/articles/2000/08/20000.html) that included the following advice:

- Don't take things personally. If someone acts inappropriately toward you, just react in a calm manner. They are likely responding that way because of outside factors.
- Admit when you don't know the answer to something. It's okay to defer to somebody else or tell the user or customer that you'll have to look into their complaint and will get back with them as soon as possible.
- It is better to validate someone's feeling or respond to the information that they have given to you than to react to them. For instance, if somebody complains that a help ticket has not been handled in a timely manner, tell them that you understand how they feel and will look into it, instead of reacting in a defensive manner.
- Don't let your personal opinions or feelings get in the way of the real complaint. Try to put yourself in the user's or customer's shoes.
- Be sympathetic. If you need a user to leave their laptop with you overnight, tell them that you realize it's frustrating and apologize.
- Try to provide a solution from which you both can benefit. Look for commonalities between you and the client, and work to find a solution that is agreeable to both of you.
- Try to be as informative as possible when discussing a solution to their problem. Most people are uncomfortable with change, so explaining the benefits of a particular solution might help ease this discomfort.
- Try to keep a positive attitude and be optimistic.
- Always work on your listening skills!

Using Appropriate Behavior

Critical to appropriate behavior is to treat the customer, or user, the way you would want to be treated. Much has been made of the Golden Rule—treating others the way you would have them treat you. Six key elements to this, from a business perspective, are punctuality, accountability, flexibility, confidentiality, respect, and privacy. The following sections discuss these elements in detail.

Punctuality

Punctuality is important and should be a part of your planning process. If you tell the customer that you will be there at 10:30 a.m., you need to make every attempt to be there at that time. If you arrive late, you have given them false hope that the problem will be solved by a set time. That can lead to anger because it can appear that you are not taking the problem seriously. Punctuality continues to be important throughout the service call and does not end with your arrival. If you need to leave to get parts and return, tell the customer when you will be back, and be there at that time. If for some reason you cannot return at the expected time, alert the customer and tell them when you can return.

Along those same lines, if a user asks how much longer the server will be down and you respond that it will up in five minutes only to have it down for five more hours, the result can be resentment and possibly anger. When estimating downtime, always allow for more time than you think you will need, just in case other problems occur. If you greatly underestimate the time, always inform the affected parties and give them a new time estimate. To use an analogy that will put it in perspective: if you take your car to get an oil change and the counter clerk tells you it will be “about 15 minutes,” the last thing you want is to be still sitting there four hours later. If you ever feel that you won’t be able to meet the timeline you proposed, communicate that as quickly as possible. It’s better to overcommunicate than to have users wondering where you are.

Exercise 10.1 is a simple drill that you can modify as needed. Its purpose is to illustrate the importance of punctuality.

EXERCISE 10.1

Understanding Punctuality

1. Consider this scenario: You call someone important in your life—your spouse, a parent, an in-law, or a close friend—and tell them that you have something very important that you need to discuss. You give that person no other details but ask them to meet you in exactly one hour at a location familiar to both of you.
2. Now imagine that you waited two hours before showing up.
3. What would be that person’s reaction? How would that person feel about having to wait for you? What kind of an impact would it have on the person’s mood and behavior?

This is an interaction with someone who matters in your life. Imagine how it would affect a customer who does not know you. Punctuality can go a long way toward keeping dialogue pleasant between two parties.

Accountability

Accountability is a trait that every technician should possess. When problems occur, you need to be accountable for them and not attempt to pass the buck to someone else. For example, suppose you are called to a site to put a larger hard drive into a server. While performing this operation, you inadvertently scrape your feet across the carpeted floor, build up energy, and zap the memory in the server. Some technicians would pretend the electrostatic discharge (ESD) never happened, put the new hard drive in, and then act completely baffled by the fact that problems unrelated to the hard drive are occurring. An accountable technician would explain to the customer exactly what happened and suggest ways of proceeding from that point—addressing and solving the problem as quickly and efficiently as possible.

Accountability also means that you do what you say you're going to do, ensure that expectations are set and met, and communicate the status with the customer. Here are some examples of ways to be accountable:

- Offer different repair or replacement options, if they're available.
- Provide proper documentation on the services that you provided.
- Follow up with the customer at a later date to ensure satisfaction.

The last one is the most overlooked, yet it can be the most important. Some technicians fix a problem and then develop an “I hope that worked and I never hear from them again” attitude. Calling your customer back (or dropping by their desk) to ensure that everything is still working right is an amazing way to build credibility and rapport quickly.

Flexibility

Flexibility is another trait that's as important as the others for a service technician. You should respond to service calls promptly and close them (solve them) as quickly as you can, but you must also be flexible. If a customer cannot have you on-site until the afternoon, you must make your best effort to work them into your schedule around the time most convenient for them. Likewise, if you are called to a site to solve a problem and the customer brings another problem to your attention while you are there, you should make every attempt to address that problem as well. Under no circumstances should you give a customer the cold shoulder or not respond to additional problems because they were not on an initial incident report.



You should always follow the express guidelines of the company you work for as they relate to flexibility, empowerment, and other issues.

It's also important that you remain flexible in dealing with challenging or difficult situations. When someone's computer has failed, they likely aren't going to be in a good mood, which can make them a “difficult customer” to deal with. In situations like these, keep in mind the following principles:

Avoid arguing or being defensive. Arguing with the customer—about anything—is only going to make the situation worse. The customer may be mad and may be yelling

at you, but don't argue back, act defensive, or take their comments personally. Try to defuse the situation by calmly reminding them that you're there to help and you want to understand what's going on so you can do that. They may need to vent for a bit, so let them do that. Just focus on doing what you need to do to resolve the problem.

Don't minimize their problems. Although the customer's problem might seem trivial to you, it isn't to them. Treat the problem as seriously as they're treating it. Keep in mind that facial expressions and body language are also important. If someone tells you their problem and you look at them like they're delusional, they're probably going to pick up on that, which can make the situation worse.

Avoid being judgmental. Don't blame or criticize. As previously mentioned, focus on what needs to happen to fix the problem. Accusing the user of causing the problem does not build rapport. Even seemingly innocent statements such as "What did you do to the computer?" can be considered judgmental and put the customer on the offensive. Stick with more neutral language, such as, "Can you help me understand what happened?"

Focus on your communication skills. If you have a difficult customer, treat it as an opportunity to see how good a communicator you really are. (Maybe your next job will be a foreign ambassador.) Ask nonconfrontational, open-ended questions. "When was the last time it worked?" is more helpful than "Did it work yesterday?" or "Did you break it this morning?" These can help you narrow down the scope of the problem.

Another good tactic here is to restate the issue or question to verify that you understand. Starting with "I understand that the problem is . . ." and then repeating what the customer said can show empathy and proves that you were listening. If you have it wrong, it's also a good opportunity to let your customer correct you so that you're on track to solve the right problem.

Confidentiality

The goal of confidentiality is to prevent or minimize unauthorized access to files and folders and disclosure of data and information. In many instances, laws and regulations require confidentiality for specific information. For example, Social Security records, payroll and employee records, medical records, and corporate information are high-value assets. This information could create liability issues or embarrassment if it fell into the wrong hands.

Over the last few years, there have been a number of cases in which bank account and credit card numbers were published on the Internet. The loss of confidence by consumers due to these types of breaches of confidentiality can far exceed the actual monetary losses from the misuse of this information.



Confidentiality entails ensuring that data expected to remain private is seen only by those who should see it. Confidentiality may be implemented through authentication and access controls.

As a computer professional, you are expected to uphold a high level of confidentiality. Should a user approach you with a sensitive issue—telling you their password, asking for assistance obtaining access to medical forms, and so on—it is your obligation as a part of your job to make certain that information goes no further.



Confidential materials on workspaces and printers should always be protected.

As part of confidentiality, don't ever disclose work-related experiences via social media. You might have had a terrible day and really want to say something like, "Wow, the people at XYZ Company sure are insufferable morons," but just don't do it. It's not professional, and it could expose you to legal action.

Respect

Much of the discussion in this chapter is focused on respecting the customer as an individual. However, you must also respect the tangibles that are important to the customer. While you may look at a monitor that they are using as an outdated piece of equipment that should be scrapped, the business owners may see it as a gift from their children when they first started their business.

Treat the customer's property as if it had value, and you will win their respect. Their property includes the system you are working on (laptop/desktop computer, monitor, peripherals, and the like) as well as other items associated with their business. Avoid using the customer's equipment, such as telephones or printers, unless it is associated with the problem you've been summoned to fix.

Another way to show respect is to focus on the task at hand and avoid distractions. For example, you should avoid the following:

- Personal calls
- Texting or social media sites
- Talking to coworkers while interacting with customers
- Personal interruptions

As for texting or talking to coworkers, there may be times when it's appropriate for you to do so based on the situation. The key is to find the right time to do it and, if appropriate, tell the customer what you are doing. For example, after gathering information, you might say something like, "Do you mind if I give my coworker Jen a quick call? The other day she told me about a situation she had that sounded exactly like this, and I want to see if her fix worked well." But then make the call quick and business-focused.



The American Customer Satisfaction Index (<https://theacsi.org>) measures the behavior of corporations and the respect they give to customers.

Respecting the customer is not rocket science. All you need to do—for this exam and in the real world—is think of how you would want someone to treat you. Exercise 10.2 explores this topic further. Like Exercise 10.1, it can be modified to fit your purpose or constraints. Its goal is to illustrate the positive power of the unexpected.

EXERCISE 10.2

Surprise Someone

1. Pick a random, toll-free number used for business solicitation and call it.
2. Chat with the operator for a few moments about the company's product or service, and then ask to speak to the supervisor.
3. When the supervisor comes on, commend the operator you were talking to for the job that they have done.

It is likely the operator became confused when you asked to speak to the supervisor; this almost always occurs only in a negative situation. How did the operator handle the request? Did it change the tone of the communication that was taking place? Did they fulfill your request even though they feared they could lose from it? Did the supervisor respond by expecting negative comments? How was the positive information you offered accepted?

Ideally, this illustrated the importance of staying professional and keeping the channel of communication open even in a tough situation. You should be able to adapt this to the workplace when a customer asks to speak to your superior or has another request that is difficult for you to fulfill.

One last area to consider that directly relates to this topic is that of ethics. *Ethics* is the application of morality to situations. While there are different schools of thought, one of the most popular areas of study is known as *normative ethics*, which focuses on what is normal or practical (right versus wrong and so on). Regardless of religion, culture, and other influences, there are generally accepted beliefs that some things are wrong (stealing, murder, and the like) and some things are right (e.g., the Golden Rule). You should always attempt to be ethical in everything you do, because it reflects not only on your character but also on your employer.

Privacy

Although there is some overlap between confidentiality and privacy, privacy is an area of computing that is becoming considerably more regulated. As a computing professional, you must stay current with applicable laws because you're often one of the primary agents expected to ensure compliance.

Although the laws provide a minimal level of privacy, you should go out of your way to respect the privacy of your users beyond what the law establishes. If you discover information about a user that you should not be privy to, you should not share it with anyone, and you should alert the customer that their data is accessible and encourage them—if applicable—to remedy the situation. This includes information that you see on their computer, on their desk, on printers, or anywhere else in their facility.



Real World Scenario

A Little Goes a Long Way

The following examples of respecting and disrespecting the customer come from one of the authors' own experiences:

"My wife and I were in an unfamiliar part of Chicago without ready access to a vehicle when we started to get hungry. I am a meat-and-potatoes man, and I rarely take a chance on anything else. There were no restaurants of that type around, however, and we wound up at an Asian grill. Expecting not to like the buffet, we ordered a side of lettuce wraps and then two buffets and drinks. As it turned out, I liked the buffet a great deal, and I went back through the line many times. We also liked the drinks and got several of those.

"Everything was great, except the waiter forgot to bring the lettuce wraps. I dismissed it and made a mental note to inform the waiter when he brought the bill and have him deduct them from our tab. Instead, the manager brought the bill over when we were finished eating, and he had scribbled on it 'no charge.' When I asked him why, he apologized that no one brought the wraps and said he hoped we would come back another time. I was beside myself with disbelief and thanked him profusely, and since then I have told many people about that restaurant, describing it as the best place in Chicago I know of to eat.

"In a very different situation, while driving home one night, the 'low tire pressure' dashboard light came on. Upon inspection, I could hear the right-rear tire hissing. I drove to a tire store and explained the situation. I had used this same tire store over the past 14 years for tires, oil changes, exhaust, maintenance, and a number of other things on the vehicles I've owned. The manager came out and said that they found a nail in the tire. They removed the nail, patched the tire, and charged me \$13. I was delighted, expecting it to cost much more, and so I paid the bill and went on my way.

"The next morning, I woke up to find the right-rear tire completely flat. I canceled the morning's appointment, filled the tire with an air compressor, and drove back to the tire store. The manager came out and told me that they found another nail in that tire; they were going to eat the \$13 on this one, but it had better not happen again. I could not believe the insinuation—that I was driving about looking for nails to hit with that one tire just so I could spend my morning taking them for \$13! Instead of offering the possibility that they had

overlooked a nail the previous night, apologizing for the inconvenience, or anything of that sort, he shifted the responsibility to me. Needless to say, I have not been back since, and all of my repair business is now done elsewhere.”

These two examples illustrate two different approaches to treating the customer. In the first example, the customer is well respected and treated better than expected. In the second example, the customer is disrespected and is treated as an inconvenience. Given the life-time value of customers, it is always better to respect them—and retain them—than to dismiss them offhandedly.

Putting It All in Perspective

Whether you are dealing with customers in person or on the phone, there are rules to which you should adhere. These were implied and discussed in the previous sections, but you must understand them and remember them for the exam.

- Use standard language and avoid using jargon, abbreviations, and acronyms. Every field has its own language, and outsiders feel lost when they start hearing it. Put yourself in the position of someone not in the field and explain what is going on by using words they understand.
- Maintain a positive attitude and tone of voice, and project confidence. The customer is counting on you to fix their problem. The last thing they want is for you to sound defeated when you hear about the problem.

Listen to your customers and take notes. Allow them to complete their statements and avoid interrupting them. People like to know that they are being heard, and as simple an act as it is, this can make all of the difference in making them feel at ease with your work.

Everyone has been in a situation where they have not been able to explain their problem fully without being interrupted or ignored. It is not enjoyable in a social setting, and it is intolerable in a business setting.

- Be culturally sensitive. Some people may have a language barrier that makes it difficult to explain their problem. (Think about how much computer language you learned in your high school language courses.) Others may have different habits or practices in their workplace. Be respectful of their world. In some cases, using the appropriate professional titles is a sign of respect, and not using them is an insult.
- Be on time. If you’re going to be late, be sure to contact your customer. Not doing so indicates that you think their problem isn’t important.
- Avoid distraction and/or interruptions when talking with customers. You need to make them feel that their problem is important and that it has your full attention. Distractions can include personal calls, texting or social media, talking to co-workers, and other personal interruptions.

- Exercise patience with difficult customers and situations.
- Avoid arguing with customers and/or becoming defensive.
- Do not minimize a customer's problem. While it may be a situation you see every day, it is a crisis to them.
- Avoid being judgmental and/or insulting or calling the customer names.
- Clarify the customer's statements and ask pertinent questions. The questions you ask should help to guide you toward isolating the problem and identifying possible solutions. Don't be afraid to nod, ask questions, and repeat to the customer what you think they are saying to make sure that you understand it correctly.
- Don't vent about customers on social media.

Set and meet—or exceed—expectations and communicate timelines and status. Customers want to know what is going on. They want to know that you understand the problem and can deal with it. Being honest and direct is almost always appreciated.

Deal appropriately with confidential materials. Don't look at files or printouts that you have no business looking at. Make sure the customer's confidential materials stay that way.

Summary

In this chapter, we covered ticketing systems and how they are implemented for an organization. We also covered the various documentation types that you will encounter as a technician throughout your career. Some of the documentation is created from external regulations. This, in turn, creates policies, which dictate processes and procedures. You'll create some of the documentation as you are troubleshooting a problem and some of the documentation after fixing a problem. We also covered documentation often used with the change management process.

Next, we looked at disaster prevention and recovery. The two main areas are data loss and equipment failure due to power issues. Data loss can be prevented with data backups and other user-facing strategies, such as Volume Shadow Copy. Power problems can be prevented with the appropriate use of uninterruptable power supplies and surge protection equipment.

Finally, we moved on to professionalism and communication. You should treat your customers as you would want to be treated. Your actions and behavior should let them know that you respect them and their business.

Exam Essentials

Understand the elements of a ticketing system. You should have a clear understanding of the importance of clear communications in the support process, as well as the important elements in a ticket, such as user information, device information, category of problem, severity of the problem, and most important, the description of the problem.

Know the common documentation in a network. Incident documentation is created when a computer, network, or security incident has happened. You should be familiar with various elements of this documentation. The acceptable use policy (AUP) is a soft policy that protects an organization's resources from employee abuse. The password policy is also a soft policy that describes best practices for password creation and expectations of the end user.

Know the elements of change management documentation. Change management is adopted by organizations to standardize the methods and procedures for any changes in the organization. Change management documentation is an integral part of this process. It is used to document a change so that it does not impede the organization's operations.

Understand how to prevent disasters and recover from data loss. Data loss is the most common disaster an organization can suffer from on a day-to-day basis. You should understand how to prevent data loss with the implementation of data backups and other prevention methods. You should also understand how to protect an organization from power-related problems using uninterruptable power supplies and surge protection.

Know the various backup methods to prevent data loss. There are three types of backups: full, incremental, and differential. Full backups copy all the data in a data set to media and reset the archive bit. An incremental backup copies only the data to media that has the archive bit set; it then resets all of the archive bits. A differential backup copies all the data in the data set to media that has the archive bit set, but it does not reset the archive bit.

Know proper communication techniques. Listen to your customers. Let them tell you what they understand the problem to be, and then interpret the problem and see if you can get them to agree to what you are hearing them say. Treat your customers with respect, whether end users or colleagues, and take their issues and problems seriously. Use professionalism when working with customers, such as minimizing distractions, speaking in plain English, being punctual, and following up.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. What is the first step to consider when entering a ticket manually?
 - A. Determine the urgency of the issue.
 - B. Collect user information.
 - C. Identify recent changes.
 - D. Verify the issue description.
2. You are executing the primary plan in the change management documentation and realize that you cannot proceed. Which section details the original configuration?
 - A. Purpose
 - B. Risk analysis
 - C. Rollback
 - D. Plan for change
3. Which section of the change management documentation contains whom the change will affect?
 - A. Business processes
 - B. Scope of change
 - C. User acceptance
 - D. Plan for change
4. You just had an outage of Internet connectivity. Which document should you complete so that stakeholders understand the reason for the outage?
 - A. Change management documentation
 - B. Knowledge base article
 - C. Acceptable use policy
 - D. Incident documentation
5. Which type of backup allows users to restore their own files without IT intervention?
 - A. Image-based backup
 - B. Volume Shadow Copy
 - C. Incremental backup
 - D. Differential backup
6. Why is it important to document the root cause in an incident document?
 - A. To satisfy legal requirements
 - B. To ensure accurate billing
 - C. To prevent the recurrence of the incident
 - D. To train new employees

7. End users are abusing the email system by selling personal items. Which policy would detail the proper use of the email system for business purposes?
 - A. MDM
 - B. Password policy
 - C. AUP
 - D. Incident management
8. Which backup media is the fastest from which to recover?
 - A. Disk-to-tape
 - B. Disk-to-disk
 - C. Disk-to-flash
 - D. Disk-to-cloud
9. You need to upgrade a server and want to make a backup of the data before you begin. Which backup method should you choose so that your normal backups are not affected?
 - A. Full
 - B. Copy
 - C. Incremental
 - D. Differential
10. What is the primary benefit of having a knowledge base in an IT environment?
 - A. Reduces the number of help desk staff needed
 - B. Ensures compliance with company policies
 - C. Provides a quick reference for resolving common issues
 - D. Improves employee morale
11. You promised a customer that you would be out to service their problem before the end of the day but have been tied up at another site. As it now becomes apparent that you will not be able to make it, what should you do? (Choose the best answer.)
 - A. Arrive first thing in the morning.
 - B. Wait until after hours and then leave a message that you were there.
 - C. Call the customer and inform them of the situation.
 - D. Email the customer to let them know that you will be late.
12. A user reports that a workstation has two significant problems that do not seem related. How should you approach these problems?
 - A. Look for what the two problems would have in common.
 - B. Assume that a virus is involved.
 - C. Deal with each issue separately.
 - D. Order a new machine.

- 13.** A customer is trying to explain a problem with their system. Unfortunately, the customer has such a thick accent that you are unable to understand their problem. What should you do? (Choose the best answer.)
- A.** Just start working on the system and look for obvious errors.
 - B.** Call your supervisor.
 - C.** Ask that another technician be sent in your place.
 - D.** Apologize and find another user or manager who can help you translate.
- 14.** You have been trying to troubleshoot a user's system all day when it suddenly becomes clear that the data is irretrievably lost. When you inform the customer, they become so angry that they shove you against a wall. What should you do?
- A.** Shove the user back, only a little harder than they shoved you.
 - B.** Shove the user back, only a little easier than they shoved you.
 - C.** Try to calm the user down.
 - D.** Yell for everyone in the area to come quickly.
- 15.** A customer tells you that a technician from your company spent three hours on the phone making personal calls. What should you do with this information?
- A.** Nothing.
 - B.** Inform your manager.
 - C.** Talk to the technician personally.
 - D.** Ask the customer to prove it.
- 16.** You arrive at the site of a failed server to find the vice president nervously pacing and worrying about lost data. What should you do?
- A.** Offer a joke to lighten things up.
 - B.** Downplay the situation and tell him that customers lose data every day.
 - C.** Keep your head down and keep looking at manuals to let him know that you are serious.
 - D.** Inform him that you've dealt with similar situations and will let him know what needs to be done as soon as possible.
- 17.** You are temporarily filling in on phone support when a caller tells you that they are sick and tired of being bounced from one hold queue to another. They want their problem fixed, and they want it fixed now. What should you do?
- A.** Inform them up front that you are only filling in temporarily and won't be of much help.
 - B.** Transfer them to another technician who handles phone calls more often.
 - C.** Try to solve their problem without putting them on hold or transferring them elsewhere.
 - D.** Suggest that they call back at another time when you are not there.

- 18.** At the end of the day, you finish a job only to find that the user you were doing it for had to leave. What should you do? (Choose two.)
- A.** Clean up and leave no evidence that you were there.
 - B.** Leave a note for the user detailing what was done and how to contact you.
 - C.** Notify the user's manager and your own manager that you have finished.
 - D.** Put the system back to its original state.
- 19.** A user on the phone does not seem to be able to explain their problem to you without using profanity. That profanity is making you unable to understand their problem. What should you do? (Choose the best answer.)
- A.** Ask the user to refrain from the offensive language.
 - B.** Overlook the profanity.
 - C.** Hang up.
 - D.** Show them that you know just as many expletives as they do.
- 20.** Which of the following is not a benefit of implementing asset tags for inventory management?
- A.** Tracking of the equipment
 - B.** Scheduling the depreciation of the equipment
 - C.** Identifying assets
 - D.** Providing ownership of the equipment

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

A user has called in and explained they accidentally overwrote a file and need to retrieve the freshest copy of the file. Luckily, you have Volume Shadow Copy configured on the share where the file was overwritten. What are the steps to recover the file?

Appendix

A

Answers to Review Questions



Chapter 1: Operating System Basics

1. B, C. You can open a command prompt by typing `cmd` or `command` in the Start menu. The command prompt utility will pop up in the search results. `Run` is not a command; it is a dialog box. `Open` is not a command; it is an operating system action.
2. D. A driver is specifically written to instruct the operating system to communicate with a piece of hardware. Source code is the original code the software is written in. An application is a program written to perform a specific purpose to interact with the user. The kernel is the core of the operating system.
3. B. OneDrive is a cloud-based file synchronization application; it is similar to Google Drive. Excel is a spreadsheet application that is part of the Microsoft 365 suite of products. Microsoft Exchange is an on-premises email server; its cloud-based counterpart is Microsoft Exchange Online, and it is similar to Google Gmail. Outlook is the email client used with Microsoft Exchange.
4. A. If you wanted to install a program on a virtual machine, the preferred method would be to mount an ISO image of the application. USB and optical discs are not common installation methods for virtual machines. Although ZIP files could be used, a ZIP file is only a compression method and not a delivery method.
5. B, C. The Windows File Explorer program can be used to copy and move files. Windows File Explorer can also be used to change file attributes. Windows File Explorer cannot be used to browse the Internet or to create backup jobs.
6. D. The maximum allowable length for a filename is 255 characters.
7. B. The shell is a program that runs on top of the OS and allows user interaction with the operating system. The Taskbar is an element of the Desktop. The Desktop does not allow for commands to be executed. Source is not related to the graphical interface.
8. B. Typing the program name in the Start box allows you to launch the program. The Shut Down command will shut down the operating system. `Run` is not a command; it is a dialog. The `cmd` command starts the command-prompt application.
9. D. When a program is multithreaded, it is written to allow for multiple requests into the processor at one time. Multiuser mode, dystopia, and preemption are not correct answers.
10. B. The Recycle Bin is a temporary spot that files are deleted to but from where they can be retrieved if deleted mistakenly. The My Computer icon is another older way to start the File Explorer. Control Panel and the Settings app are used to customize the operating system.
11. D. In Windows 11, when you click the Start button, a power icon is displayed on the lower left. From the power icon, you can sleep, reboot, or shut down the operating system. There is no command named `turnoff`. Turning the power switch off and unplugging the machine are not graceful ways to power off the operating system. Pressing `Ctrl+Alt+Del` alone will not power down the computer.

- 12. A. A Trusted Platform Module (TPM) 2.0 is required for the installation of Windows 11. 64 GB or greater of storage is required for the installation of Windows 11. A Microsoft DirectX 12 graphics device is required for the installation of Windows 11. The minimum amount of memory required for the installation of the 64-bit Windows 11 operating system is 4 GB. More RAM is always recommended, but only 4 GB is required.
- 13. D. The minimum amount of free hard drive space required for Windows 11 64-bit is 64 GB.
- 14. D. The minimum amount of memory required for the 64-bit installation of the Windows 11 is 4 GB.
- 15. C. A quick method of accessing Help is to press the F1 key. The Windows key on the keyboard launches the Start menu. The F12 key and the Alt key perform various tasks, depending on the application.
- 16. A. The Windows key+R will open a run box, where you can launch applications by typing in their executable name. Ctrl+R is used to reload web pages in web browsers. The Windows key alone will launch the Start menu. The F1 key will launch help in applications and on the desktop it will launch a web browser.
- 17. C. The notification area is the area on the rightmost portion of the Taskbar. It is used to display the date and time and running background applications. The pinned apps and Start menus are located on the center-left portion of the Taskbar. Shutdown options are found inside the Start menu in Windows 11.
- 18. A. Clicking the Settings gear in the Start menu will open the Settings app. From there, the System category will allow you to change the Display Properties. Clicking the System icon under Control Panel allows you to change system settings. Pressing Ctrl+Alt+Esc launches the Task Manager. Pressing Ctrl+Alt+Tab launches the Task View in Windows 11.
- 19. B. The File Explorer allows you to navigate the file/folder structure in Windows 11. The Start menu allows you to launch applications. KDE and GNOME are Linux desktop shells.
- 20. C. The minimum processor speed required for the 64-bit installation of the Windows 11 is 1GHz.

Chapter 2: Windows Configuration

- 1. B. Task Scheduler (`taskschd.msc`) will allow you to start a program based on a condition met, such as the computer starting up. The Programs and Features applet is not an MMC snap-in. The Disk Management MMC snap-in allows you to manage the disks attached to the computer. The Group Policy Editor lets you edit the local Group Policy.
- 2. C. FAT32 is a filesystem that was popular with Windows 9x operating systems and is also commonly used for USB flash drives for compatibility between platforms. NTFS is the native filesystem for Windows. EFS is a technology that encrypts files. The CDFS filesystem is used for CD media.

3. C. The shutdown utility can be used to schedule a remote shutdown—for example, `shutdown /t 60 /m \\computer`. The `taskmgr` utility is used to view tasks, `kill` is used to kill processes, and `netstat` is used to view network statistics and activity.
4. D. The System event log would contain an entry when the operating system reboots. The Application log is used to log applications installed on the system. The boot log is not a valid log. The Security log contains log information pertaining to the security of objects and files.
5. D. The exFAT filesystem is a Microsoft proprietary filesystem created for large flash drives. GPT is a partition table type, NFS is a network filesystem used with Linux/UNIX, and ext3 is a filesystem used with Linux.
6. D. The Accounts settings screen allows you to access the Backup screen, where you can configure backups for the user under the Windows Backup menu item. The System section allows you to configure a multitude of options for Windows itself. The Devices section lets you configure devices, and the Apps section allows you to uninstall and configure applications.
7. C. The Apple File System (APFS) is a proprietary filesystem created for SSD support; it can be found on macOS Sierra 10.12.4, iOS 10.3, and tvOS 10.2 and later operating systems. Hierarchical File System (HFS) is a proprietary filesystem that was developed by Apple for the Mac OS back in 1985, ext4 (Fourth Extended File System) is used with Android and other Linux-based operating systems, and exFAT is a Microsoft proprietary filesystem created for large flash drives.
8. A. The System Control Panel applet is used to change the operating system computer name and change its membership from a workgroup to a domain, or domain to workgroup. The User Accounts applet is used to make changes to your user account or others on the operating system. The Windows Tools option in Control Panel opens up a menu of tools, none of which join the operating system to a domain. The Programs And Features applet is used to uninstall and repair installed programs, as well as manage Windows features.
9. D. When a program is not responding to any commands, the process must be terminated via the Task Manager. You can end the process by opening the Task Manager, right-clicking the application, and selecting End Task. Viewing the performance will not make the process respond. Adding more memory is not possible without shutting down the computer. Pressing Ctrl+Alt+Delete without any further action will not reboot the computer.
10. B. The command `eventvwr.msc` will start the Event Viewer snap-in. The command `eventviewer.exe` is not a valid command. The command `lusrmgr.msc` will start the Local Users and Group snap-in. The command `devmgmt.msc` will start the Device Manager snap-in.
11. C. The iSCSI Initiator utility can be found in Control Panel, in Windows Tools. The Storage Spaces applet is used to configure storage spaces of already connected storage. The Disk Management snap-in is used to configure disks already connected to the operating system. Device Manager is used to configure hardware attached to the computer.
12. A. Hardware resources are configured in the Device Manager console. Files and folders and applications are configured in the operating system GUI. Memory is configured in System Properties.

13. C. The mechanical hard drive should be periodically defragmented to arrange concurrent blocks of storage of files so that the hard drive arm does not need to move much when reading and writing files. Solid-state drives (SSDs) do not suffer poor performance from fragmentation, because they have no moving parts. Freeing up space and emptying the Recycle Bin are good ideas but do not contribute to the hard drive's performance as much as defragmenting. Trim is a feature used with SSD drives, where the operating system notifies the SSD device that data blocks have been freed up.
14. A. The Indexing Options applet allows you to configure custom extensions so that metadata can be searched. The File Explorer Options applet is used to view and configure the File Explorer. The Windows Tools option in Control Panel is not really an applet; it opens up a menu of tools. The Ease of Access Center applet is used to configure options related to user disabilities to make the operating system easier to use.
15. A. The operating system must boot from an active primary partition if MBR style partitioning is used. Extended partitions are used to contain logical partitions. There is no such thing as a dynamic partition.
16. B. The HKEY_LOCAL_MACHINE Registry hive contains information about the computer's hardware. It is also known as HKLM. HKEY_CURRENT_MACHINE and HKEY_MACHINE are not valid Registry hives. HKEY_RESOURCES was used with Windows 9x operating systems but is no longer used.
17. C. GUID [globally unique identifier] Partition Table (GPT) allocates a 64-bit logical block addressing entry to partition tables to accommodate larger drives. GPT has rapidly replaced the MBR-style partition tables. LILO and GRUB are boot managers for Linux. NFS is a network filesystem used with Linux/UNIX operating systems.
18. A. Disk Defragmenter will rearrange blocks of files on a hard drive to create contiguous files and space. Windows File Explorer is used to navigate the filesystem. Scandisk is not a command. Windows Backup allows for file backups.
19. A. The hibernate power mode will write the contents of the RAM to the disk so that the computer can enter a low power state. The standby power mode will turn off the monitor and lower power consumption by the computer. The sleep and suspend power mode will lower power consumption, but it will not write the contents of the RAM to the disk.
20. C. The Dynamic Disk partition style allows up to 2,000 volumes to be created on a physical disk. GUID [globally unique identifier] Partition Table (GPT) allow up to 128 volumes to be created on a physical disk. The Master Boot Record (MBR) and Basic Disk partitioning scheme are one and the same and allow up to 26 partitions to be created on a physical disk.

Chapter 3: Windows Administration

1. D. Windows 11 requires a minimum of 4 GB of RAM and 64 GB of storage. Therefore, to meet these requirements, you need to upgrade both the RAM and the storage. Simply proceeding with the installation would not work because the current specifications do not meet

the minimum requirements. Upgrading only the RAM or only the storage would still leave one of the requirements unmet, making both upgrades necessary.

2. A. The Generalize pass is where the operating system detects hardware and installs the appropriate drivers. The Out-Of-Box Experience (OOBE) is the pass that is responsible for configuring and creating the user environment during the setup process. The Specialization pass is where the operating system is configured during setup. The WinPE pass starts the procedure of installing the operating system.
3. D. The WinRE partition is where the Windows Recovery Environment is located. The BCD is a term used with the boot configuration data to direct the operating system how to boot. The system partition contains the kernel of the operating system. The EFI system partition (ESP) contains the boot configuration data (BCD) that is used to boot the rest of the operating system on a UEFI system.
4. A. The Windows Deployment Services (WDS) is a server role that is used to create a Windows image for mass deployment. The Microsoft Assessment and Planning (MAP) toolkit is used to create reports on system hardware for image deployment. The User State Migration Toolkit (USMT) is used in a corporate environment to migrate user data. The `sysprep` utility is used to prepare an operating system for imaging.
5. D. The `robocopy` command will copy all data and will include NTFS permission to remain intact. The `xcopy` and `copy` commands will copy files from a source folder to a destination folder but will not copy NTFS permissions. The `chkdsk` command is used to check the integrity of the NTFS filesystem.
6. B. The General Availability branch will install updates as they are released to the general public. The Semi-Annual Channel is no longer a channel supported by Windows 11. The Long-Term Servicing Channel never installs new features during the life of the version of Windows. The Insider Program allows for the installation of brand-new features before they are publicly released.
7. A. Windows 11 Home only allows for the General Availability branch releases. Windows 11 Pro allows for the use of the Insider Program branch releases. Windows 11 Education allows for the use of the Insider Program branch releases. Windows 11 Enterprise allows for the use of Insider Program branch releases.
8. B. Unified Extensible Firmware Interface (UEFI) firmware is required for Secure Boot, which enhances security by ensuring that only trusted software is loaded during the boot process. Basic Input/Output System (BIOS) and Legacy BIOS do not support Secure Boot. Extensible Firmware Interface (EFI) is an older term; UEFI is the correct modern term.
9. C. The `pathping` command will measure the packet loss at each router as the packet travels to the destination address; it combines the `ping` and `tracert` commands. The `ping` command will return a single destination's response time. The `nslookup` command is used to resolve DNS addresses. The `tracert` command allows you to see how a packet travels to its destination.

10. B. The `msinfo32.exe` tool allows for the remote reporting of a computer's hardware. `regedit.exe` is used to edit the Registry. `msconfig.exe` is used to change the startup of services and change the boot process. `dxdiag.exe` is used to diagnose DirectX problems.
11. C. The `chkdsk` command is used to check a volume for corruption, as well as attempt to repair the corruption. The `diskpart` command allows you to create, modify, and view volumes on a disk. The `format` command allows you to format a filesystem on a volume. The `sfc` command is used to fix corrupted files but not volume corruption.
12. C. Users will not be able to change the wallpaper or personalize the operating system until it is active. Applications will still launch fine. Windows Updates will continue to function to keep Windows up-to-date. Browsing the Internet will not be affected.
13. B. The Out-Of-Box Experience (OOBE) is the pass that is responsible for configuring and creating the user environment during the setup process. The Generalization pass is where the operating system detects hardware. The Specialization pass is where the operating system is configured during setup. The WinPE pass starts the process of installing the operating system.
14. A. When installing Windows, you can control telemetry data in the Privacy Settings dialog box. The Locale options allow you to control how the language and local notations of money, weight, date, and several other measurements for a particular region of the world. Disk partitioning is configured during the initial installation in order to direct Windows where to install. Account creation is performed to create the first user on the operating system.
15. B. The `tracert` command shows the path packets take to a destination, helping to identify where delays or failures occur. The `ping` command checks connectivity but not the path. `nslookup` resolves DNS names but does not trace the path. `netstat` shows network connections and ports but not the path.
16. D. There is no way to directly convert a 32-bit installation of Windows to 64-bit Windows without performing a clean installation. Upgrading directly to Windows 11 Pro 64-bit is not possible. Upgrading to Windows 11 Pro 32-bit is possible, but you cannot upgrade to a 64-bit edition. Upgrading to Windows 11 64-bit is not possible.
17. A. The BCD is stored in the EFI System Partition on an EFI installation of Windows. The WinRE partition is used for the Windows Recovery Environment. Secure Boot is a feature of an EFI installation and does not contain its own partition. The `C:\WINDOWS` folder is where the installation of Windows exists.
18. A. The `netstat` command can be used to view ports in use by the operating system and network applications communicating with the network. The `ipconfig` command allows you to see the current IP address and DNS information for the operating system. `pathping` allows you to view packet loss along the path to a destination IP address. The `nslookup` command is used to resolve DNS records.
19. C. The Wireless Wide Area Network (WWAN) connection is a cellular connection and is configured on the Cellular screen in the Settings app. A wired connection is configured on the

Ethernet screen in the Settings app. A wireless connection is configured on the Wi-Fi screen in the Settings app. A virtual private network (VPN) connection is configured on the VPN screen in the Settings app.

20. C. The maximum number of connections that can simultaneously be made to a Windows workstation is 20. All other answers are incorrect.

Chapter 4: Working with macOS and Linux

1. C. The command `ls -la` will list all the files in a long format. The command `ls -a | ls -l` will not work. `ls -s`; `ls -a` will show two listings: one with the size and the other with all the files. The `ls -a\ls -l` command will show two listings—one with all the files and the other in a long format—but it will not show all the files in a long format.
2. B. Mission Control is a quick way on macOS to view what is currently running. Keychain or keychain access is a tool that saves credentials so that the user does not need to be prompted. The Finder is the macOS equivalent to the Windows File Explorer. Force Quit is a function of the macOS for killing a process.
3. A. Time Machine is the Apple application that can be used to create backups on a regular basis. The Finder is the macOS equivalent to Windows File Explorer. VSS is a Windows technology that is used to create snapshots of a volume. Keychain is a tool that saves credentials so that the user does not need to be prompted.
4. B. The `nano` command is used to edit files. The `ps` command lists processes running. The `rm` command removes files or directories. The `ls` command lists files and folders in the filesystem.
5. C. The `ip` command can be used to edit an Ethernet connection's configuration settings. The `dd` command is used to duplicate disks. The `apt` command is used with the APT package management system for downloading packages. The `pwd` command will show the current working directory.
6. A. The Passwords feature has replaced the Keychain tool in the latest version of macOS 15.2 (Sequoia). The Passwords feature saves credentials so that the user does not need to be prompted. Spotlight is a utility in macOS that allows a search of the operating system and Internet for search terms. The Dock is the macOS application menu that is located on the bottom of the screen.
7. A. The shell is the interpreter in Linux between the operating system and the user. The translator is a wrong answer. Login is something you perform to access the shell. The graphical user interface (GUI) does not interpret commands and therefore is not considered an interpreter.

8. D. Working copies are backups that are kept on-site at the computer center for immediate recovery purposes. An on-path attack is an attack where an eavesdropper listens between two people or devices. Cloud copies are backups that are sent to cloud storage and cannot be immediately recovered. Journal copies is a wrong answer.
9. C. The `apt` utility can be used to download and apply patches to a Linux installation. The `update` command is not a utility. Shell/terminal is an interface for interacting with the operating system with the command line. The `patch` command is not a utility.
10. C. The `chown` command is used to change ownership of a file. The `cd` command changes the working directory. The `chmod` command changes permissions on files. The `pwd` command displays the current working directory.
11. A. The `fsck` Linux utility is used to check and repair disks. The `chkdsk` utility is a Windows utility used to check and repair disks. The `du` utility is used to show the current disk usage. `dumgr` is not a utility and is a wrong answer.
12. B. By pressing the Home button twice on an iPad, you can Force Quit an application that will not stop running. The `kill` utility can be used only at the command line of Linux/macOS. The Task Manager is a Windows utility. Close Quit is not a feature and therefore a wrong answer.
13. C. Bash is the most common command-line shell used with Linux. Tcl/Tk is a shell used with Linux, but it is very uncommon. Terminal is a utility used to access the command-line shell. SSH, or Secure Shell, is a secure protocol used to access the shell over a network.
14. D. The Dock is the area at the bottom of the macOS screen that is used to launch applications. The footer is not a valid answer. Mission Control is a macOS feature that allows you to see all the applications running. The Taskbar is a Windows Desktop element that is used to track running applications and quickly launch them.
15. C. The Command+spacebar key combination will bring up the Spotlight utility. Control+Shift, Option+Tab, and Alt+Home are wrong answers.
16. A. The `sudo` command can be used to run a single command as another user. The `su` command allows you to change user logins at the command line. The `passwd` command changes the user's password. The `ifconfig` command allows you to view and modify the wired network interface.
17. C. The `ps` command will display a snapshot of the current running processes on a Linux operating system. The `ls` command will display a listing of files from the working directory. The `cat` command will display the contents of a file. The `su` command allows you to change user logins at the command line.
18. B. The command `cd ..` will take you one level back from the current working directory. The command `cd .` will do nothing, because the period signifies the current working directory. The command `cd . . .` is not a valid command. The command `cd ~` will change directories to the home directory of the user.

19. B. The effective permissions are read and write. From left to right, the permissions are `rwX` for the user, `rw-` for the group, and `r--` for everyone else. Since the user is only a member of the group applied to the file, they will have read and write permissions.
20. C. The `-p` option on the `mkdir` command allows subfolders to be created as well as the target folder. All other answers are incorrect.

Chapter 5: Security Concepts

1. A. Fences are intended to delay or deter entrance into a facility. Access control vestibules are used for mid-layer access control to prevent tailgating. Multifactor authentication is used for mid- and inner-layer access control. Strong passwords are used for mid- and inner-layer access control.
2. B. A magnetometer, also known as a metal detector, can detect weapons on a person entering a facility. Biometrics are used to authenticate someone based on their face, retina, fingerprint, or some other method that biologically verifies their identity. A motion sensor is used to detect motion. A badge reader is used to scan security ID badges.
3. A. Social engineering uses the inherent trust in the human species, as opposed to technology, to gain access to your environment. IDSs are network-based systems that detect intrusions. Perimeter security describes physical security. Biometrics describes an authentication method based on human physical traits.
4. C. A fingerprint scanner, or any device that identifies a person by a physical trait, is considered a biometric security control. Access control is the system that controls access for users. Physical barriers are structures that limit physical access. Softening refers to weakening of security.
5. A. Although the end result of any of these attacks may be denying authorized users access to network resources, a denial-of-service (DoS) attack is specifically intended to prevent access to network resources by overwhelming or flooding a service or network. Worms reproduce and move throughout the network to infect other systems. Trojans are programs that enter a system or network under the guise of another program. Social engineering is a process in which an attacker attempts to acquire information about your network and system by social means, such as talking to people in the organization.
6. B. A distributed denial-of-service (DDoS) attack uses multiple computer systems to attack a server or host in the network. A denial-of-service (DoS) is a one-on-one attack to disrupt service. Worms reproduce and move throughout the network to infect other systems and therefore do not attack one victim. Rootkits are software programs that have the ability to hide themselves from the operating system.
7. D. Tokens are rotating numerical keys that you must physically have with you to gain access to the operating system or applications. Biometrics relies on a physical characteristic of the user to verify identity. Biometric devices typically use either a hand pattern or a retinal scan

to accomplish this. Smartcards contain a private certificate key and are protected with a password. Geofencing uses your GPS coordinates to ensure that the authentication happens when you are in a defined geographic area.

8. A. An on-path attack intercepts data and then sends the information to the server as if nothing were wrong, while collecting the information. Zero-day attacks are attacks in which a developer has not properly patched a hole yet and is unaware of the hole. A denial-of-service (DoS) attack is used to disrupt legitimate requests from being answered. A Smurf attack is a type of distributed denial-of-service (DDoS).
9. A. A spoofing attack is an attempt by someone or something to masquerade as someone else (IP address) and is often used to disrupt access. Social engineering is a process in which an attacker attempts to acquire information about your network and system by social means, such as talking to people in the organization. Worms reproduce and move throughout the network to infect other systems. Password attacks are used in an attempt to guess passwords.
10. A. A worm is different from a virus in that it can reproduce itself, is self-contained, and doesn't need a host application to be transported. A smurf attack is a type of distributed denial-of-service (DDoS). A phishing attack is an attempt to gain a user's credentials to a network resource. Trojan horses are programs that enter a system or network under the guise of another program.
11. D. The reflective attack is using a broadcast ping (ICMP) on a network. The return address of the ping may be that of a valid system in your network. The Transmission Control Protocol (TCP) is not typically used with a reflective attack. The Internet Protocol (IP) is a suite of protocols and solely used with a reflective attack. The User Datagram Protocol (UDP) is not described in this reflective attack.
12. C. A SQL injection attack is a method of passing a SQL query with a web request by using an escape code sequence. An insider threat is a threat from within your organization, such as a disgruntled employee. An evil twin attack involves a rogue access point with the same SSID as your organization. Tailgating is the act of walking behind someone who has swiped to get into an area so the attacker can gain entry.
13. B. A key fob is an example of authentication for something you have. A password is something you know. A fingerprint is something you are. A place is a geographical place in which you are.
14. B. End-user education is the best way to protect your users from the threat of phishing via email. Antivirus software is used to prevent viruses, not phishing attempts. SecureDNS can be useful in protecting your users, but not from phishing emails. The principle of least privilege assigns only the permissions that users need to do their work, and no more.
15. C. Spear phishing is a type of social engineering, where someone is trying to con your organization into revealing account and password information by pretending to be a high-level person. A spoofing attack is an attempt by someone or something to masquerade as someone else, with the intent of disrupting access. A replay attack is a form of an on-path attack, where packets are replayed at a critical time. Trojan horses are programs that enter a system or network under the guise of another program.

16. A. Some viruses won't damage a system in an attempt to spread into all the other systems in a network. These viruses use that system as the carrier of the virus. A password attack would not prompt your antivirus software to notify you. Your antivirus software could be malfunctioning, but it would not suggest the same virus is infecting you over and over again. A denial-of-service (DoS) attack would not prompt your antivirus to notify you.
17. B. A symptom of many viruses is unusual activity on the system disk. The virus spreading to other files on your system causes this. A disk failure will not create high disk activity. A spear phishing attack is a social engineering attack and will not create high disk activity; neither denial-of-service attacks nor TCP/IP hijacking attacks will create high disk activity.
18. A. A brute-force attack is a type of password attack in which a password is guessed over and over, until the right password is guessed. A backdoor attack is an embedded account that allows unauthorized access through an unpatched coding hole. A worm is different from a virus in that it can reproduce itself, is self-contained, and doesn't need a host application to be transported. A TCP/IP hijacking attempt is an attack that attempts to redirect the TCP/IP conversation to the threat agent.
19. A. The thesis of your presentation should outline the need of a firewall to isolate the external network from the internal network. Firewalls will not scan packets for viruses. Firewalls will not prevent password attacks or harden physical security.
20. C. A login script is used by Active Directory during login to map drives and printers. A home folder is a private network location in which the user can store their personal files. Organizational units (OUs) are used to group computers and users so that Group Policy can be applied. The MMC is used to manage various aspects of Active Directory and the local operating system.

Chapter 6: Securing Operating Systems

1. C. A bring your own device (BYOD) policy defines the minimum specifications for an employee's device used for work-related access. The mobile device management (MDM) software would usually police these specifications, but it would not define them. The acceptable use policy (AUP) is a code of conduct when dealing with organization resources. The nondisclosure agreement (NDA) is an agreement used when dealing with intellectual property.
2. D. Synchronizing a mobile device with a computer system allows you to mirror personal data between the devices, regardless of which one contains the most current data. Calibration refers to matching the device's and the user's perceptions of where the user is touching the screen. Remote wipes allow you to remove personal data from a lost or stolen device. Pairing is what must be done in Bluetooth for two Bluetooth devices to connect and communicate.
3. A. The Guest account should be disabled on the operating system, unless there is good reason to leave the account enabled. Print Operators is a group found on Windows servers. Power Users is a group found on both Windows workstations and servers. Userone is

obviously a user account. Unless the user has left the organization, there should be no reason to disable the account.

4. C. Locator apps can find the phone and then display a message that can be read even while the phone is locked. Locator apps also often give the user the option to make the phone blare a noise to aid in finding the device. Failed login restrictions will wipe a device if a specific number of wrong passwords are entered. Antivirus apps make sure that your phone does not get a virus, but they offer no way to message or locate the device. Remotely wiping a mobile device is a feature of locator apps, but it is not itself an app that displays messages and produces noises.
5. C. A screensaver should automatically start after a short period of idle time, and that screensaver should require a password before the user can begin the session again. The system should neither shut down nor restart, since work could still be open. For the same reason, the system should not log out the user.
6. C. User Account Control (UAC) helps protect against unauthorized changes to system files by prompting for administrative approval. BitLocker and EFS provide encryption. Windows Defender offers antivirus protection.
7. C. The goal of single sign-on (SSO) is to allow a username/password combination. Once the combination is entered, claims are used to access additional resources. Tokens are given to the operating system after a user successfully logs in; they allow a user to access rights on the operating system. Kerberos is used (along with Active Directory) to authenticate a user on the Windows operating system. Multifactor authentication is the use of two or more factors to authenticate a user.
8. B. The Windows Credential Manager is a password manager that is built into the operating system of Windows. Edge and Internet Explorer 11 work in conjunction with the Credential Manager, but they are not password managers. Active Directory is used to authenticate domain users, but it does not manage passwords for end users.
9. A. The advantage to assigning the IP addresses statically is that you can make certain which host is associated with which IP address, and then use filtering to limit network access to only those hosts. Disabling DNS will limit all hosts on a network from getting to the Internet and will not limit network access. A default subnet mask will not limit network access and will cause other problems. An empty default gateway will not limit network access.
10. C. Wi-Fi Protected Access 2 (WPA2) offers the Advanced Encryption Standard (AES) for encrypting wireless communications. Wired Equivalent Privacy (WEP) offers weak 64- or 128-bit encryption. Wi-Fi Protected Access (WPA) uses the RC4 encryption algorithm. Temporal Key Integrity Protocol (TKIP) is a part of the WPA encryption protocol.
11. B. An extension is a type of add-on that will extend the functionality of the web browser in a way it was never originally designed. A pop-up blocker changed the way a web page is rendered. A plug-in changes the way a web page is rendered. An ad blocker also changes the way a web page is rendered.

12. A. When an employee is offboarded, their user account is deleted or disabled. A user account is not created, nor is its password reset, during the offboarding process. A user account normally is created during the onboarding process.
13. B. By default, a Google account is required on Android devices; it is used to synchronize data and app purchases to the cloud. The device does not require email to be set up, but the account can be used for the setup. The account is not used for desktop backups. The device does not require registration.
14. A. Fingerprint locks are the most secure of all the lock methods, since fingerprints are hard to duplicate. Face locks have a high number of false positives, which can be used to gain access to the phone. Passcode locks can be cracked or shouldered. Swipe locks are not really locks; they just allow the screen to be opened with a swipe.
15. A. The Encrypted File System (EFS) is a functionality of the Windows NTFS filesystem. EFS can encrypt individual files and folders. BitLocker is a full-device encryption technology. NTFS is a filesystem that supports encryption and security, among other functionality. BitLocker to Go is used for full-device encryption of removable drives.
16. B. Because the user is in both groups and the Sales group has full share permissions and the Sales group has modify NTFS permissions, the most restrictive of the two is Modify, so that will be the effective permission for the user. All of the other answers are incorrect.
17. A. The permissions will be the same as before the move, since you are just moving the files and not creating a new entity. The permissions will not be inherited from the parent folder. The permissions will not be configured the same as the root folder. The permissions will not be blank.
18. B. The user will have only the Modify permission when logged in locally to the computer, since the filesystem is not shared with the appropriate permissions. The user will not have the Modify permission when connecting from the network. The user will still have the Modify permission when logged in locally, because of the NTFS permissions. The user will not have read-only permissions when connecting from the network.
19. C. The file attributes are most likely set to Read-only, and thus you are not permitted to delete the file. The Modify permission would not hinder the file from being deleted. The share permissions would not have any bearing on the problem, since the local filesystem is being accessed. The file attributes being set to System would not restrict a file from being deleted.
20. D. Mobile device management (MDM) software enables you to enforce profile security requirements on mobile devices. The acceptable use policy (AUP) is a code of ethics your users should follow when dealing with organizational resources. A nondisclosure agreement (NDA) is an agreement between an employee and the organization to protect intellectual property. A bring your own device (BYOD) policy explains how devices should be secured but provides no enforcement.

Chapter 7: Troubleshooting Operating Systems and Security

1. D. Windows Update is responsible for downloading and installing Windows service packs, patches, and security updates. Device Manager is used to view devices installed on the operating system. The Microsoft Management Console is a console that allows snap-ins to be added for management. Download Manager is a component of Edge.
2. B. Enable Safe Boot with the `msconfig` utility allows you to boot with basic drivers and minimal startup of nonessential services. Enable Debugging is used by kernel developers. Disable Driver Signature Enforcement is used to allow an unsigned driver to load during boot. Enable Low-Resolution Video will boot the operating system into a VGA mode.
3. B. The `/REBUILDBCD` option can be used with the `bootrec` tool to rebuild the boot configuration data (BCD). The `/FIXBOOT` option writes a new boot sector to the system partition. The `/SCANOS` option scans all other partitions that are found to have Windows installations. The `/FIXMBR` writes a new master boot record (MBR) to the partition.
4. B. The most important first step is to identify and verify the malware symptoms. You should quarantine the infected system once you have verified it is infected. Remediating the infected system happens after you disable System Restore. Education of the end user is the last step to malware removal.
5. A. The Profile tab of the Advanced Systems Properties dialog box allows you to view the total size of a local or remote profile. Regedit and the Windows Recovery Environment will not aid in troubleshooting a slow-loading profile. Windows Preinstallation Environment is the mini-Windows version used for installation of Windows.
6. D. `winresume.exe` is used to load Windows from a suspended state. The Boot Configuration Data (BCD) is used to direct Windows to boot the proper installation. `ntoskrnl.exe` is the Windows kernel. `winload.efi` is used for the normal booting of the Windows operating system.
7. A. The System Restore option should be used first to restore the operating system to an earlier point before the problem. This will restore the device back to a previous state before the installation of the drivers. System Restore will not affect user data files. Reset This PC will reset the PC back to factory default before the tax application was installed. Reimaging the laptop will erase all programs and data files. Manually reinstalling Windows 10 will erase all programs and data files.
8. A. When you have UEFI firmware, you must have the disk set up with a GUID Partition Table (GPT) partitioning type. The standard master boot record (MBR) partitioning type can be used with BIOS. Power-on self-test (POST) is a routine the BIOS or firmware performs to test hardware before boot. The Boot Sector is contained on both MBR and GPT partitioning types.

9. B. Pop-up blockers are used to prevent pop-ups and pop-unders from appearing. Antimalware utilities will remove and prevent malware. Phishing sites are used to collect users' credentials by tricking users. Antivirus software is used to protect the operating system from viruses.
10. C. Antivirus definitions should be updated daily, because new viruses are identified by the minute. Updating antivirus definitions weekly or monthly will open you up to the possibility of infection.
11. A. Resource Monitor can be used to identify slow and sluggish performance, as well as identify the source of the problem. The `msconfig.exe` tool can be used to enable or disable services on startup and launch tools, but it cannot be used to diagnose performance issues. The Device Manager MMC can be used to view and modify devices, but it will not help diagnose performance problems. Reliability Monitor will display the reliability of the operating system, but it will not help diagnose problems with performance.
12. B. Windows Update Troubleshooter can assist in diagnosing problems with Windows Update. The `ntbtlog.txt` file is used to diagnose problems with bootup. Windows Recovery Environment is used to solve problems with Windows and is not typically used for problems with Windows Updates. Safe mode is a boot mode that loads minimal drivers and services.
13. A. Windows Defender Security is considered antimalware and antivirus protection for the Windows operating system. Mobile device management (MDM) software is used to manage mobile devices. Windows Action Center is a notification center for action to be taken in the operating system. VirusTotal is a third-party site that analyzes virus signatures, but it does not protect you from them.
14. B. The System File Checker (SFC) allows you to manually scan for modified operating system files and repair them. `regedit` is used to modify the Registry. `bootrec` is used to repair the boot records on an operating system installation. User Account Control (UAC) is used to control access to administrative credentials.
15. A. Establishing security policies and procedures will help eliminate security problems and guide employees on what to do if they arise. Optimizing drives will defragment drives and has no effect on security. Preventing booting into safe mode will only hinder diagnostics. Preventing booting into Windows Recovery Environment will also hinder diagnostics.
16. D. The device most likely has too many background applications open that are using RAM. That the device is not charged to capacity would not affect RAM. A digitizer not functioning properly would resemble inaccurate touchscreen responses. If the device were in Do Not Disturb (DND) mode, the speakers would not work.
17. B. If auto-reconnect is configured on an SSID, the device could be susceptible to an evil twin attack, in which the device connects to any device with the same SSID. The device will not reconnect to any SSID but only to the SSID configured as auto-reconnect. Battery life will remain unaffected with auto-reconnect. Exceeding limits can be avoided with auto-reconnect, because the wireless network is used for data usage.

18. B. Event Viewer will allow you to see more detailed information on why programs have crashed. The Event Viewer logs may not give the exact reason, but they will aid in understanding the root cause. Device Manager is used to view and manage devices connected to the operating system. The Windows Recovery Environment is used to repair the Windows operating system. The `msconfig.exe` tool is used to modify startup programs and launch other diagnostic tools.
19. C. The reason that the operating system would write out large amounts of RAM to the page file is that the system is running low on physical RAM and is attempting to free up physical RAM. The CPU might run high when paging occurs, but it will not trigger excessive paging. Although it is normal for the operating system to write out RAM to the page file, it is not normal for large amounts to be written out. The page file is not faster than conventional RAM.
20. B. Degraded battery life can be expected from an overheating mobile device if the problem persists for a long time. Higher RAM usage will not occur with overheating, but it could be a cause of overheating. Inaccurate touchscreen responses are not a symptom or a consequence of overheating. The inability to decrypt emails depends on having the proper certificate installed.

Chapter 8: Scripting, Remote Access, and Artificial Intelligence

1. C. Scripting languages are interpreted languages that run on top of a runtime environment. Programming languages, not scripting languages, require a compiler. Scripting languages are not strongly typed; programming languages are strongly typed. Scripting languages have bad memory management because of loosely typed variables.
2. A. Scripting languages are considered high-level languages because they do not directly access hardware and use an intermediary called the interpreter. Mid-level languages are Java and C/C++, not scripting languages. There is no such thing as an intermediate-level language. Low-level languages are machine language and assembly language, which are not scripting languages.
3. C. Floating-point variables allow for precision math, also known as decimal math. Boolean variables allow for true or false values. Integer variables allow for whole numbers values. String variables allow for text values.
4. C. A program variable is the least significant and it is not inherited. A system variable is defined for the entire system and is the most significant because it is inherited by all users and programs. A user variable is significant as well, since all applications will inherit the variable. A string variable is not inherited, but it is not an environment variable.

5. B. The statement `$xvar = 2` is a PowerShell statement that will load the variable `xvar` with a value of 2. The statement `xvar = 2` is Bash syntax. The statement `xvar = 2;` is JavaScript syntax. The statement `set /a xvar=2` is Windows batch script syntax.
6. D. A `for` loop has a defined beginning and end, and steps from the beginning to the end. A `do while` loop is a type of `while` loop and has only a defined end. A `while` loop has only a defined end. An `if` statement is branch logic, not a loop.
7. C. The `.bat` extension is used with the Windows batch scripting language. The `.vbs` extension is used with VBScript language. The `.js` extension is used with the JavaScript scripting language. The `.py` extension is used with the Python scripting language.
8. B. The VBScript language allows for the use of the Component Object Model (COM). The PowerShell scripting language allows for the use of the .NET Framework. Windows batch scripts use existing applications. JavaScript is primarily web browser-based and does not allow for the use of external objects.
9. D. The `.py` extension is used with the Python scripting language. The `.vbs` extension is used with the VBScript language. The `.js` extension is used with the JavaScript scripting language. The `.bat` extension is used with the Windows batch scripting language.
10. A. The PowerShell scripting language allows for the use of the .NET Framework and is commonly used with Microsoft Azure and Microsoft 365. The VBScript language allows for the use of the Component Object Model (COM) and is not used for cloud services. A Windows batch script uses existing applications and is not used for cloud services. JavaScript is primarily web browser-based, does not allow for the use of external objects, and is not used for cloud services.
11. D. JavaScript is primarily web browser-based and allows for interactive content. The PowerShell scripting language is used to manage the operating system. The Bash scripting language is primarily used with Linux and UNIX systems. Windows batch scripts use existing applications and are used to manage the operating system.
12. B. The `.sh` extension is used with the Bash scripting language. The `.vbs` extension is used with the VBScript language. The `.bat` extension is used with the Windows batch scripting language. The `.py` extension is used with the Python scripting language.
13. C. Before a Bash script can be executed, you must use the `chmod` command to grant execute permissions. The `chown` command changes ownership. There is no such thing as an `execute` attribute. Adding `.sh` to the end of the script doesn't serve any purpose.
14. C. The statement `mvar = 8;` is JavaScript syntax to load a variable of `mvar` with a value of 8. The statement `$mvar = 8` is PowerShell syntax. The statement `mvar = 8` is Bash syntax. The statement `set /a mvar=8` is Windows batch script syntax.
15. D. PowerShell is the only scripting language that has a preinstalled Integrated Scripting Environment (ISE), called the PowerShell ISE. The VBScript language requires the installation of Microsoft Visual Studio Code. The Bash scripting language requires a text editor or

other package to be installed. The Python scripting language requires a third-party integrated development environment (IDE), also known as an ISE.

16. A. The line `//comment` is used to comment JavaScript code. The line `'comment` is used to comment VBScript code. The line `REM comment` is used to comment Windows batch script code. The line `# comment` is used to comment Bash script code and PowerShell code.
17. A. The `.js` extension is used with the JavaScript scripting language. The `.sh` extension is used with the Bash scripting language. The `.bat` extension is used with the Windows batch scripting language. The `.py` extension is used with the Python scripting language.
18. B. The primary goal of AI is to simulate or imitate human intelligence. This involves processes like learning, thinking, and problem-solving. AI is not intended to replace human intelligence but to complement it. Creating emotional experiences is not the main goal of AI. AI focuses on cognitive processes, not physical abilities.
19. B. Convolutional neural networks (CNNs) are specialized for image and video analysis. Language translation is typically handled by recurrent neural networks (RNNs). Speech recognition is also handled by RNNs. Game playing can be handled by various types of neural networks.
20. A. Tokenization is the method of dividing text into smaller parts known as tokens. Tokenization does not involve converting images into text. Tokenization is not about translating text. Tokenization deals with text, not audio data. Translating text into another language is the function of transformers.

Chapter 9: Safety and Environmental Concerns

1. C. A three-pronged grabber should be used to retrieve the screw from the computer. Using a magnetic-tipped screwdriver is not advisable, because many components are sensitive to magnets. Using a magnetic grabber is not advisable, because many components are sensitive to magnets. Shaking the computer until it falls out can be dangerous and make the screw harder to get out.
2. B. Using your personal computer to manage your side business adheres to the AUP and prevents misuse of company resources. Using your work computer violates the AUP. Asking for permission is unnecessary as the AUP clearly prohibits such use.
3. B. The material safety data sheet (MSDS) will have the necessary information about the cleanup procedures. The Occupational Safety and Health Administration (OSHA) is an organization that oversees safety in the workplace. The product label may or may not have this information. A cathode-ray tube (CRT) is a type of monitor that uses a glass tube.

4. D. You can download the MSDS by visiting the printer cartridge manufacturer's website. The store is not required to furnish an MSDS at the time of purchase. The manufacturer is not required to furnish an MSDS with the packaging of the product. You are legally allowed to have an MSDS for products, as per OSHA. You should have them on hand in case of an emergency.
5. A, C. An accident should always be reported so that if there is a hazardous condition, it can be fixed. Report hazardous conditions, such a spill on the floor, to the employer before they are a problem. A near-accident does not need to be reported unless it is caused by a hazardous condition, like spills. Rain forecasted for a workday is not a hazardous condition.
6. B. You can feel an ESD shock of 3,000 volts or more. 300 volts is too low for a human to feel. A minimum of 30,000 volts can be felt by a human. 300,000 volts can be felt by humans and could be painful.
7. A, B. Using an antistatic wrist strap allows you to ground yourself to dissipate a static charge. Using antistatic bags allows for a uniform charge around a component. Spraying disinfectant spray will only kill germs, not reduce static. Shuffling your feet will build a charge of static.
8. B, D. Providing properly maintained tools and equipment is a requirement for an OSHA-compliant workplace. Displaying an OSHA poster in a prominent location is also a requirement for an OSHA-compliant workplace. Attending yearly OSHA safe work environment seminars is not a requirement. Having an OSHA employee stationed within five miles of the facility is also not a requirement.
9. A. You should ensure that the company has the proper licenses to install 20 additional copies to be compliant with licensing. Microsoft Office is not an open source licensed product. Using the personal license key from an existing system violates the license agreement. Following normal installation procedures will not ensure compliance of licensing.
10. C, D. You should use a cart to move heavy objects. You should also ensure that no safety hazards are in your path. Lifting by bending over at the waist can hurt your back. Carrying CRT monitors with the glass facing outward is not safe, as the weight will be farthest from your body.
11. C. Alkaline batteries should be taken to a recycling center. Throwing batteries in the trash is not environmentally responsible. Incinerating batteries is not advisable, since they can explode and will create pollution. Flushing batteries down the toilet is not an acceptable disposal method.
12. B, D. The first step is to confiscate and preserve the prohibited materials on the drive. The next step is to report the prohibited materials through the proper channels. Destroying the prohibited material will not remedy the situation. Confronting the user about the material is not suggested, as they could become hostile.
13. A. You should research information on the Payment Card Industry Data Security Standard (PCI DSS) standard. The General Data Protection Regulation (GDPR) is used for protecting EU citizens. Protected health information (PHI) is any data that defines a patient or an

ailment of a patient. Personally identifiable information (PII) is any information that can be used to identify a person.

14. A. Wood and paper fires can be put out by a Class A fire extinguisher. Class B fire extinguishers are used for flammable liquids. Class C fire extinguishers are used for electrical fires. Class D fire extinguishers are used for flammable metals.
15. B, C. Phillips and Torx are two common types of screwdrivers. Neither circular nor helix is a type of screwdriver.
16. A, B. Good safety plans protect the interests of the workers and also help to keep company costs down. Periodic workplace inspections and a training program are good components to implement. Employees cannot be punished in any way for reporting safety violations. Third-party audits of safety programs are not necessary but can complement periodic inspections and safety and health training programs.
17. A. A good rule of thumb when it comes to magnetic-tipped screwdrivers is to avoid using them inside a computer case. Magnetic tools can damage data on disks that use a magnetic storage scheme. Magnetic-tipped screwdrivers should not be used inside a computer even if they are kept away from the processor and the RAM. Regardless of whether they are powered or not, magnetic-tipped screwdrivers should not be used inside computers.
18. B, C. Material safety data sheets (MSDSs) contain handling and storage instructions as well as personal protection instructions. MSDSs do not contain freezing point specifications or weight.
19. B, D. To maintain a safe work environment, all employees must follow certain protocols. These include using protective gear and equipment and following all health and safety rules. Accidents must be reported to the employer, not OSHA (unless action is not properly taken by the employer). There are no rules requiring safety training, although it is a good idea to have these sessions and for employees to attend them.
20. B. Nickel-metal hydride (NiMH) batteries are not considered environmental hazards. Alkaline batteries used to contain mercury, which is an environmental hazard, so they are assumed to be environmental hazards. Nickel-cadmium (NiCd) batteries and button cell batteries are considered environmental hazards.

Chapter 10: Documentation and Professionalism

1. B. Collecting user information is correct because it is essential to know who is reporting the issue and how to contact them. Determining the urgency of the issue is important, but communication to the customer is of higher importance. Identifying recent changes would happen after you enter the ticket. Verifying the issue description is an important, but secondary, step.

2. C. The rollback section contains the original configuration that can be used to revert the changes. The purpose section contains the reason for the proposed change. The risk analysis section explains the risks involved with the proposed changes. The plan for change section contains the primary change configuration and the alternate change configuration.
3. B. The scope of change section details whom the change will affect. The business processes section details the current business processes the change will affect. The end-user acceptance section details how the changes were tested and accepted by the users. The plan for change contains the primary and alternate plans for the proposed change.
4. D. Incident documentation should be completed so that key stakeholders can understand the reason for the outage. Change management documentation is used for proposed changes to the network. Knowledge base articles are used to document symptoms and solutions. An acceptable use policy (AUP) is used to protect an organization's resources from user abuse.
5. B. Volume Shadow Copy allows users to restore previous versions of their files themselves. Image-based restoration typically requires a desktop administrator. Incremental and differential backup and restores also require IT intervention for restoration.
6. C. Documenting the root cause helps in preventing the incident from recurring. There are typically no legal requirements around the documentation of the root cause, unless there is something specifically in a contract. Billing is also not related to documenting root cause, unless there is something specifically in a contract. Training new employees is important but is not a primary reason for documenting the root cause.
7. C. The acceptable use policy (AUP) details the acceptable use of the email system for business purposes. Mobile device management (MDM) is software that allows you to manage mobile devices in the workplace. A password policy details the appropriate handling and management of passwords. Incident management is how a network or security incident is handled.
8. B. Disk-to-disk is the fastest recovery method and backup method as well, because you are backing up from a disk to another disk attached via the network. Disk-to-tape is slower because you must re-tension the tape and then locate the data on the tape to recover it. Disk-to-flash is not a backup method, because of the price of flash. Disk-to-cloud is the slowest recovery method because you must recover from the cloud over a network connection.
9. B. You should use the copy backup method, since it will perform a full backup of the files without resetting the archive bits. A full backup makes a full backup and resets all the archive bits affecting the normal backups. An incremental backup copies only the files that have changed since the last backup and leaves the archive bits unchanged. A differential backup backs up only the files that have changed since the last backup and then resets all the archive bits.
10. C. Providing a quick reference for resolving common issues helps both users and technicians solve problems more efficiently. Creating a knowledge base rarely ever reduces help desk staff. Knowledge bases do not ensure compliance. Although knowledge bases are great tools, they do not improve morale.

11. C. Calling the customer and informing them of the situation is the best action that can be performed, since you are having direct communications. Arriving first thing in the morning should be done, but communication is the first action. Waiting until after hours and then leaving a message is not an appropriate action. Sending an email letting the customer know that you will be late is lying, since you will not make it to their site.
12. C. If the problems do not appear to be related, then deal with them separately. Looking for commonalities between the problems would just waste valuable time. No assumptions should be made, such as that a virus is causing the problem, unless there is proof. Ordering a new machine is disruptive to the customer, unless the problem dictates that this action be performed.
13. D. If you do not understand the customer, you should apologize, treat the customer with respect, and seek a manager to help translate. Ignoring the customer and starting to work on the system is not an appropriate response. Calling your supervisor will not solve the problem, since your supervisor will have the same issue. Asking that another technician be sent in your place is being insensitive to the customer and “passing the buck.”
14. C. You should avoid confrontations with the user by attempting to calm them down. You should never shove the user back. Yelling for everyone in the area to come quickly will escalate the situation, which is the opposite of what you want to do.
15. B. You should take appropriate action and inform your manager, as well as notify the customer that you’ve done so. Doing nothing is not an appropriate response to this situation. Talking to the technician personally will create a confrontation and should be avoided. Asking the customer to prove it is an inappropriate action to this problem.
16. D. You should assure the vice president that you are optimistic and skilled to deal with these problems. Offering a joke is an inappropriate action. Downplaying the situation does not show respect to the customer’s problem. Keeping your head down and ignoring the customer does not display appropriate communications.
17. C. You should try to solve the customer’s problem without further escalating the frustration by putting them on hold or transferring them. Informing them that you won’t be able to help is only going to anger the caller, further validating their frustration. Transferring them to another technician is not the appropriate action, because it will add to their frustration. Suggesting that they call back will create more frustration and is poor customer service.
18. B, C. Leaving a note for the user detailing what was done and how to contact you displays appropriate communications. Notifying the user’s manager and your own manager that the problem is resolved is also appropriate communications. Cleaning up and leaving no evidence that you were there is an inappropriate action, as it is deceiving to the user. Putting the system back to its original state is an inappropriate action because it does not solve the problem.
19. A. While the user’s profanity is likely linked to frustration, it hinders the communication and should be eliminated. You should ask the user to refrain from the offensive language. There may be circumstances where it is necessary to overlook the profanity, if your request is likely

to make the already upset customer even angrier, as long as you can understand what is being said. Hanging up is not an appropriate action and will further the user's frustration. Firing back with profanity is not an appropriate response, as it will escalate the situation further.

- 20.** B. Scheduling of the depreciation of the equipment is performed in accounting software. Tracking of the equipment is a benefit to an asset tag. Identifying assets is a direct benefit of asset tags. An asset tag provides a proof of ownership.

Appendix

B

Answers to Performance-Based Questions



Chapter 1: Operating System Basics

Answer to Performance-Based Question


In order to accommodate the future requirement of BranchCache, your organization will need to purchase a volume license agreement with Microsoft. The BranchCache feature is only available in Windows 11 Enterprise. Windows 10 Pro is a retail operating system that can be upgraded to Windows 11 Enterprise. However, the upgrade will require a different 25-digit product key and activation of the Windows 11 Enterprise operating system.

Chapter 2: Windows Configuration

Answer to Performance-Based Question

To check to see which background processes are running and the resources they are using, open Task Manager. Do so by pressing Ctrl+Alt+Delete and selecting Task Manager. You can also press Ctrl+Shift+Esc.

Once in Task Manager, click the Processes tab, as shown here. If all the processes are not shown, then expand the More Details chevron in the lower left, if it is not already expanded. Click the CPU column header to sort by CPU usage. If a process is taking up a considerable amount of CPU time, you can highlight it and click End Process to shut it down. You can also sort by memory used and shut down processes that look to be using excessive amounts of memory. Note that shutting down critical processes may cause Windows to lock up or otherwise not work properly, so be careful what you choose to terminate. Note also that right-clicking one of the processes offers the End Process Tree option—a useful option when the process being killed is associated with others.




Chapter 3: Windows Administration

Answer to Performance-Based Question


- Assuming you are on a version of Windows 11, the process is as follows:
1. Click the Start menu.
 2. Click Settings.
 3. Select Accounts.
 4. Click Access Work Or School.
 5. Click Connect.
 6. When the Set Up A Work Or School Account dialog box appears, click Join This Device To A Local Active Directory Domain as the alternate action, as shown in Figure 3.61.

FIGURE 3.61 The Set Up A Work Or School Account dialog box




7. Enter the domain name in the dialog box, as shown in Figure 3.62.

FIGURE 3.62 Domain dialog box




8. Enter the user account on the domain that has rights to join workstations, as shown in Figure 3.63.

FIGURE 3.63 Domain administrative credentials dialog box



9. Enter the domain user and their respective security level on this PC, as shown in Figure 3.64. Or you can click Skip, and the user will automatically have the same rights as the domain.

FIGURE 3.64 Local administrative rights dialog box




10. Click the Restart Now option and reboot.

Chapter 4: Working with macOS and Linux

Answer to Performance-Based Question


The listing you see when typing these commands will differ based on such factors as the system, the directory, your permissions, and the files/subdirectories present, but in all cases, there will be entries present with the `-a` option that do not appear in the display without it. Among those listings that now appear are a single period (representing the present directory) and a double period (representing the parent directory), as shown here:



```

user@localhost:~$
[user@localhost ~]$ cd /
[user@localhost /]$ ls -F
afs/  boot/  etc/   lib@  media/  opt/   root/  sbin@  sys/   usr/
bin@  dev/   home/  lib64@  mnt/   proc/  run/   srv/   tmp/   var/
[user@localhost /]$ ls -aF
./  afs/  boot/  etc/   lib@  media/  opt/   root/  sbin@  sys/   usr/
../ bin@  dev/   home/  lib64@  mnt/   proc/  run/   srv/   tmp/   var/
[user@localhost /]$ cd ~
[user@localhost ~]$ ls -l
total 4
-rw-r--r--. 1 user user 216 Jan  1 22:24 file.html
drwxr-xr-x. 2 user user  6 Jan  1 20:24 images
[user@localhost ~]$ ls -al
total 24
drwx-----. 4 user user  145 Jan  1 22:24 .
drwxr-xr-x. 5 root root   48 Jan  2 23:22 ..
-rw-----. 1 user user 3285 Jan  3 00:56 .bash_history
-rw-r--r--. 1 user user   18 Aug 11 20:00 .bash_logout
-rw-r--r--. 1 user user  144 Aug 11 20:00 .bash_profile
-rw-r--r--. 1 user user  522 Aug 11 20:00 .bashrc
drwx-----. 3 user user   20 Jan  1 12:47 .config
-rw-r--r--. 1 user user  216 Jan  1 22:24 file.html
drwxr-xr-x. 2 user user   6 Jan  1 20:24 images
-rw-----. 1 user user   20 Jan  1 16:26 .lessht
[user@localhost ~]$
  
```

If there are any files or directories starting with a period, they will now appear where they did not before. The easiest way to “hide” a file or directory in Linux is to start the name of it with a period; thus, it will not show up in a listing unless the `-a` option is used. An example of this is shown here:



```

user@localhost:~$
[user@localhost ~]$
[user@localhost ~]$ ls -l
total 4
-rw-r--r--. 1 user user 216 Jan  1 22:24 file.html
drwxr-xr-x. 2 user user  6 Jan  1 20:24 images
[user@localhost ~]$ mv file.html .file.html
[user@localhost ~]$ ls -l
total 0
drwxr-xr-x. 2 user user 6 Jan  1 20:24 images
[user@localhost ~]$ ls -al
total 24
drwx-----. 4 user user 146 Jan  3 22:51 .
drwxr-xr-x. 5 root root  48 Jan  2 23:22 ..
-rw-----. 1 user user 3285 Jan  3 00:56 .bash_history
-rw-r--r--. 1 user user  18 Aug 11 20:00 .bash_logout
-rw-r--r--. 1 user user 144 Aug 11 20:00 .bash_profile
-rw-r--r--. 1 user user 522 Aug 11 20:00 .bashrc
drwx-----. 3 user user  20 Jan  1 12:47 .config
-rw-r--r--. 1 user user 216 Jan  1 22:24 .file.html
drwxr-xr-x. 2 user user  6 Jan  1 20:24 images
-rw-----. 1 user user  20 Jan  1 16:26 .lessht
[user@localhost ~]$

```

Chapter 5: Security Concepts

Answer to Performance-Based Question

A simple 8-character alpha numeric password contains 0–9 for a total of 10 characters, 26 uppercase and 26 lowercase characters. This gives you a total of 52 letters and 10 numbers, for a total of 62 combinations per character: 62 to the power of 8, or $62 \times 62 \times 62 \times 62 \times 62 \times 62 \times 62 \times 62 = 218,340,105,584,896$ combinations. A 25-character alphanumeric password with symbols contains 95 combinations per character; 95 to the power of 25 is 2.77×10^{54} combinations. If you are using a calculator, you might see $2.7738957e+49$ as a result. Although the exact math is not significant, the deep understanding of combinations and complexity is the underlying lesson.

Chapter 6: Securing Operating Systems

Answer to Performance-Based Question

The following explains how you would achieve the goal:

1. Create three folders on the local filesystem of the server: Sales, Marketing, and R&D.
2. Set NTFS permissions on all folders for Everyone and the Modify permission.

3. Set the Sales folder to have the following share permissions:
 - **Sales:** Change
 - **Marketing:** Read
4. Set the Marketing folder to have the following share permissions:
 - **Marketing:** Change
 - **Sales:** Read
5. Set the R&D folder to have the following share permissions:
 - **R&D:** Change
 - **Marketing:** Read

Chapter 7: Troubleshooting Operating Systems and Security

Answer to Performance-Based Question

1. Investigate and verify malware symptoms.
2. Quarantine the infected systems.
3. Disable System Restore in Windows.
4. Remediate infected systems.
 - a. Update antimalware software.
 - b. Scan the removal techniques, such as booting into Safe Mode and using the Windows Preinstallation Environment.
 - c. Reimage/reinstall.
5. Schedule scans and run updates.
6. Enable System Restore and create a restore point in Windows.
7. Educate the end user.

Chapter 8: Scripting, Remote Access, and Artificial Intelligence

Answer to Performance-Based Question

To write a PowerShell script to find other scripts in a user profile directory and all its subdirectories, you need the \$home environment variable. There are a number of ways of writing this script to achieve the solution. The following is just one of the possibilities, using the parameters found on the Microsoft website:

```
Get-ChildItem -Path $home\* -Include *.ps1 -Recurse
```



PowerShell cmdlets and their associated parameters are well documented on Microsoft's website. For more information about the Get-ChildItem cmdlet, visit:

<https://learn.microsoft.com/en-us/powershell/module/microsoft.powershell.management/get-childitem?view=powershell-7.5>

To learn about other PowerShell cmdlets, visit:

<https://learn.microsoft.com/en-us/powershell/scripting/overview?view=powershell-7.5>

Chapter 9: Safety and Environmental Concerns

Answer to Performance-Based Question


Here are some steps to take to look for trip hazards and eliminate them:

1. Walk around the server room and count how many cables are lying on the floor.
2. Walk around the client areas and see how many cables are lying on the floor or are exposed underneath cubicles.
3. Devise a plan to secure the cables and prevent them from being hazards. For example, purchase and install floor guards, use *cable ties*, or install conduits, as necessary, to secure all loose cables.


Chapter 10: Documentation and Professionalism

Answer to Performance-Based Question

1. Browse to the folder containing the overwritten file.
2. Right-click the file and select Properties.
3. Select the Previous Versions tab.
4. Find the latest version of the file and click the Restore button, as shown here:



5. In the confirmation box, select Restore again, as shown here:



Index

A

- About section (System section), 145
- acceptable use policy (AUP), 729–730
- access control, 389, 392
- access control lists (ACLs), 380–382
- access control vestibule, 366
- access denied, 584
- Access Work Or School section (Accounts screen), 152
- Accessibility Settings screen (Settings app), 153–155
- Accessibility tab (System Settings screen), 286–288
- Account Protection screen (Settings app), 132–133, 135
- accountability, 674, 788
- accounts, 434–439
- Accounts section (Settings app), 151–152
- accuracy, artificial intelligence (AI) and, 673
- Activation section (System section), 145
- active detection, 724–725
- Active Directory (AD), 22, 31, 243–244, 389–396
- Active Directory Users and Computers (ADUC), 391
- active window, in windows, 55
- Activity History section (Privacy Settings screen), 141
- ad blockers, 500
- administration (Windows)
 - about, 176, 260
 - command-line tools
 - about, 217
 - commands with standard privileges *vs.* administrative privileges, 235
 - disk commands, 230–232
 - informational tools, 232–234
 - navigation and file management, 218–222
 - network connectivity tools, 223–227
 - OS management tools, 227–230
- exam essentials, 260–261
- exercises, 219–220, 232
- installing Windows
 - about, 176–179
 - boot options, 212–216
 - considerations, 216–217
 - options for, 179–180
 - process for, 180–200
 - recovery partition, 206
 - repair installation, 205–206
- networking
 - about, 236
 - accessing resources, 249–252
 - client network configuration, 256–259
 - establishing connections, 244–249
 - firewall settings, 252–255
 - models for, 236–242
 - user authentication, 243–244
- performance-based question answers, 827–830
- performance-based questions, 266
- review question answers, 805–808
- review questions, 262–265
- upgrading Windows
 - about, 176–179
 - boot options, 212–216
 - considerations, 216–217
 - editions, 209
 - image deployment, 207
 - process for, 200–205
 - security and feature updates, 209–212

- side-by-side upgrading Windows
 - 10/11, 206–207
 - zero-touch deployment, 207–209
- administrative privileges, commands
 - with, 235
- Administrative shares, 252, 474–475
- Administrative Tools and Windows Tools
 - applet, 95, 125–127
- Administrator account, 459, 465–466
- ads, number of, 613
- Advanced App Settings section (Apps Settings screen), 140
- Advanced Configuration and Power Interface (ACPI), 113
- Advanced Encryption Standard (AES), 441–442
- Advanced Network Settings section (Network & Internet section), 149
- Advanced Package Tool (APT), 304, 316
- Advanced Power Management (APM), 113
- advanced power savings, 113
- Advanced RISC Machine, 10
- Advanced tab (System applet), 104–109
- adware, 398, 404–405
- air filter masks, 699–700
- Airplane Mode, 603, 605
- alarm system, 369
- algorithmic bias, 670
- alkaline batteries, 719
- alpha stage, in operating system life cycle, 14
- alternate click, 43
- American Customer Satisfaction Index, 790
- amplified DoS, 418, 419
- Android, 12
- Android package (APK), 609
- Android phones, setting passcode
 - locks on, 522
- antenna and access point placement, in SOHO networks (wireless), 511
- anti-malware, 406–407
- antistatic bags, 692
- antistatic foam, 693
- antistatic mats, 691
- antistatic wrist straps, 689–691, 699
- antivirus, 405–406
 - antivirus engine, 406
 - antivirus/antimalware updates, 305
- App & Browser Control screen (Settings app), 133–134
- APP filenames, 274
- App Permissions section (Privacy Settings screen), 143
- App Store, installing applications
 - from, 272–273
- appearance, 781–782
- Apple filesystem (APFS), 159–160
- Apple ID, 273
- Apple Macintosh OS (macOS), 11
- applets, Control Panel, 95–98
- application spoofing, 611
- applications
 - about, 7, 17–18
 - allowing, 482–483
 - authenticator, 529
 - bootleg, 611
 - closing, 603
 - cloud-based productivity tools, 23–32
 - considerations for, 23
 - crashes, 553–557, 598–599
 - installing, 20–22, 644–645
 - on macOS, 272–276
 - malicious, 611
 - mobile OS problems, 597–600
 - requirements and resources, 570
 - security considerations, 22–23
 - system requirements for, 18–20
 - updating, 529, 603
- /Applications folder, 296
- appropriate use, for artificial intelligence (AI), 674–675
- Apps for Websites section (Apps Settings screen), 141
- Apps Settings screen
 - Advanced App Settings section, 140
 - Apps for Websites section, 141
 - Default Apps section, 140

- Installed Apps section, 139–140
- Offline Maps section, 141
- Settings app, 139–141
- Startup section, 141
- Video Playback section, 141
- architecture, requirements for, 18
- Archive attribute, 471
- arguments, 307
- armored virus, 401–402
- ARP poisoning, 421, 422
- ARP spoofing, 421, 422
- arrays, adding, 164
- artificial intelligence (AI)
 - about, 623, 660–661, 675
 - components
 - considerations, 673–675
 - datasets, 667–669
 - foundation models, 665–667
 - integration, 669–670
 - limitations, 670–673
 - machine learning (ML), 662–665
 - natural language processing (NLP), 665
 - evolution, 661–662
 - exam essentials, 675–676
 - exercises, 669–670
 - performance-based question answers, 834
 - performance-based questions, 680
 - review question answers, 817–819
 - review questions, 677–679
- assembly language, 623
- Assessment and Deployment Kit (ADK), 207
- asset management, 755–758
- asset tags, 756
- assigned users, 758
- asterisk (*), 59, 221, 330
- at sign (@), 330
- attributes, for files and folders, 471–474
- Audio section (Accessibility Settings screen), 155
- audio transformers, 666
- audits, 389, 392
- authentication, 382, 388, 512–513
- authentication, authorization, and accounting (AAA), 513
- authenticator applications, 529

- authorization, 388
- autodiscover record, 28
- automated backups, scripting and, 645
- automated installation, 22
- Automatic Sample Submission, 481
- automation, scripting and, 640–642
- AutoPilot, 188–189, 209
- AutoPlay feature, 440
- AutoPlay section (Bluetooth & Devices Settings screen), 148
- autorotate issues, for mobile OS, 607–608
- AutoRun, disabling, 439–440
- Azure Active Directory (Azure AD), 31–32

B

- background processes, limiting, 603
- background refresh cycle, 227
- Background section
 - Display Settings window, 44
 - Settings app, 136
- backout plan, 765–766
- backup and recovery
 - about, 288, 299–302, 773
 - automated, 645
 - backup media, 777–778
 - backup methods, 779
 - backup plan, 766
 - backup rule, 778
 - backup testing, 780–781
 - critical applications, 777
 - data restoration, 780
 - file-level backups, 773–776
 - image-based backups, 776–777
- badge reader, 367
- barcode, 756
- bare-metal backup, 776–777
- basic disks, 162–163
- Basic Input/Output System (BIOS), 179, 431, 548
- Basic Input/Output System/Unified Extensible Firmware Interface (BIOS/UEFI), 19
- batteries, 603–604, 719–720
- behavior, appropriate, 788–793

best practices

for change management

about, 765

change board, 772

change type and schedule, 769–770

documented business

processes, 765–766

end-user acceptance, 772–773

purpose of change, 767–768

request forms, 767

risk analysis, 770–772

scope of change, 768–769

for incident response, 728–729

for macOS

about, 299

antivirus/antimalware updates, 305

driver/firmware updates, 305

patch management, 304

Rapid Security Response

(RSR), 305–306

scheduled backups, 299–302

scheduled disk maintenance, 302

system updates/App Store, 302–303

for malware removal, 589–597

for security

about, 428–429

account management, 434–439

data encryption, 429

disabling AutoRun, 439–440

end-user, 432–434

password expiration, 432

requiring passwords, 431

setting strong passwords, 429–430

beta stage, in operating system life cycle, 14

bias, 670–671, 672

biometrics, 376–378, 463

BitLocker, 37–38, 121, 486

BitLocker to Go, 486–487

Blue Screen of Death (BSOD), 557–558

Bluetooth & Devices Settings screen (Settings app), 145–148

bollards, 373–374

Boot Configuration Data (BCD), 179, 213, 549, 550

Boot Manager, 550

boot process, 212–216, 324–325, 548–552

boot sector virus, 402

Boot tab (System Configuration tool), 88, 89

bootleg applications, 611

border, in windows, 55

botnets, 404

Bourne shell, 636–637

branch logic, 629

brownouts, 707, 710

browser-related, 585–589

brute-force attacks, 422–423

built-in video cards, requirements for, 19

business email compromise (NEC), 425

Business Productivity Online Suite (BPOS), 25

button cell, 719, 720

BYOD devices, as a vulnerability, 428

BYOD policy, 530–531

byte-pair encoding (BPE), 665

C

cable locks, 370–371

cable ties, 701

cache, clearing, 497–498

Cached Exchange Mode, 27–28

Calculator, 55

Cameras section (Bluetooth & Devices Settings screen), 147

campus license, 732

cans, 722

Captions section (Accessibility Settings screen), 155

captive portal, 246

Captures section (Gaming Settings screen), 149

Cascading Style Sheets (CSS), 639–640

cases, 687

cat command, 337

categories, for operating systems, 10–11

cathode-ray tube (CRT), 686–687

cd command, 218–219, 328–329

cd images command, 342

- central processing unit (CPU), 18, 543
- certificate authority (CA), 487, 494–495
- certificate issues, 586–587
- Certificate Manager snap-in, 83–84
- certificate of destruction, 444
- certificate of recycling, 444
- chain of custody, 726
- change advisory board, 772
- change board, 772
- change freeze, 770
- change management
 - best practices
 - about, 765
 - change board, 772
 - change type and schedule, 769–770
 - documented business
 - processes, 765–766
 - end-user acceptance, 772–773
 - purpose of change, 767–768
 - request forms, 767
 - risk analysis, 770–772
 - scope of change, 768–769
 - plan document, 765
- change type and schedule, 769–770
- Change Type option (Advanced tab of System applet), 107–108
- character tokens, 665
- chemical solvents, 722
- chkdsk command, 231–232
- chmod command, 637
- chown command, 333
- ChromeOS, 12
- Chromium, 504
- cleaning systems, 711–713
- clearing
 - browsing data, 496–497
 - cache, 497–498
- clicks, in Windows, 43
- client network configuration, 256–259
- clientless VPN, 653
- client-to-site VPN, 652–653
- Clipboard section (System section), 145
- Close button, in windows, 54
- closed source, 8
- closed-circuit television (CCTV), 367
- closing apps, 603
- cloud licensing, 32
- cloud-based deployment, 208–209
- cloud-based Email security
 - gateways, 408–409
- cloud-based operating systems, 11
- cloud-based productivity tools, 23–32
- Cloud-Delivered Protection, 481
- cmdlets, 631
- collaboration tools, 29–31
- collapsing folders, 59
- Color Filter section (Accessibility Settings screen), 154
- Color section (Display Settings window), 44
- Colors section (Settings app), 136
- Command Prompt, 629
- command-line tools
 - about, 217
 - commands with standard privileges *vs.* administrative privileges, 235
 - disk commands, 230–232
 - informational tools, 232–234
 - navigation and file management, 218–222
 - network connectivity tools, 223–227
 - OS management tools, 227–230
- commands, 309–312. *See also specific commands*
- comment syntax, 627
- common documentation, 758–763
- Common Language Runtime (CLR), 631
- communication, clarity of, 754
- Compact Disc File System (CDFS), 161
- companion virus, 402
- compliance, 730–731
- Component Object Model (COM), 635
- components
 - artificial intelligence (AI)
 - considerations, 673–675
 - datasets, 667–669
 - foundation models, 665–667
 - integration, 669–670

- limitations, 670–673
- machine learning (ML), 662–665
- natural language processing (NLP), 665
- macOS
 - about, 276
 - features, 289–295
 - system settings, 279–288
 - tools, 277–279
- compressed air, 699
- compression options, 472–473
- Computer Management snap-in, 78–79
- Computer Name tab (System applet), 104
- computer vacuum, 712
- computers, 685–688, 713–722
- concepts, for operating systems, 8–10
- concurrent license, 732
- conditional access, 386
- confidentiality, 789–790
- configuration
 - about, 72, 166–167
 - alternative IP addresses in
 - Windows, 258–259
 - Disk Management, 158–165
 - exam essentials, 167
 - exercises, 77, 86, 98, 103
 - interacting with operating systems
 - about, 72
 - Control Panel, 95–127
 - Disk Cleanup, 93–95
 - Resource Monitor, 92–93
 - System Configuration tool, 88–90
 - System Information tool, 90–92
 - Task Manager, 72–77
 - Windows Settings app, 127–155
 - performance-based question, 171
 - performance-based question
 - answers, 826–827
 - real world scenarios, 157–158
 - review question answers, 803–805
 - review questions, 168–171
 - Windows Registry, 155–158
- configuration management database (CMDB), 756
- configuration profiles, 525
- confirmation bias, 671
- connectivity issues, for mobile OS, 604–607
- considerations, for installing and
 - upgrading, 216–217
- content filtering, 515–516, 524
- Continuity, 293–294
- continuous monitoring and
 - improvement, 674
- Contrast Themes section (Accessibility
 - Settings screen), 154
- Control box, in windows, 54
- Control Panel
 - about, 95–98
 - Administrative Tools and Windows Tools
 - applet, 125–127
 - applets, 95–98
 - BitLocker applet, 121
 - configuring network addressing, 258
 - Credential Manager applet, 116
 - Device Manager applet, 119–120
 - Devices and Printers applet, 117, 118
 - Ease of Access Center applet, 125, 126
 - File Explorer Options applet, 100–103
 - Indexing Options applet, 123–125
 - Internet Options applet, 99–101
 - Mail applet, 123
 - Network and Sharing Center applet,
 - 119, 120
 - Power Options applet, 111–115
 - Programs and Features applet, 116–117
 - Regional and Language Options, 99
 - Sound applet, 117, 118
 - System applet, 103–110
 - Troubleshooting applet, 117–119
 - User Accounts applet, 110–111
 - Windows Defender Firewall
 - applet, 121–122
- Control Program for Microcomputers
 - (CP/M), 217
- Controlled Folder Access, 481–482
- cooperative multitasking, 9
- copper wires, as a source of interference, 694

copy backups, 779
 copy command, 220–221
 Copy To option (Advanced tab of System applet), 108
 cordless phones, as a source of interference, 694
 corporate license, 732
 corporate-owned mobile devices, 531
 costs, of iCloud, 293
 Counter Cipher Mode with Block Chaining Message Authentication Code Protocol (CCMP), 508
 coverage, frequency and, 510
 cp command, 334
 CPU tab (Resource Monitor), 92
 creative industries, hallucinations and, 672
 Credential Manager, 116, 463, 493–494
 critical applications, 777
 cross-site scripting (XSS), 424–425
 cryptominers, 398
 curl command, 349–350
 current, voltage *vs.*, 686
 current directory, 308–309, 328
 current stage, in operating system life cycle, 14
 customer communication, 782–786
 customer service issues, hallucinations and, 672
 Cybersecurity and Infrastructure Security Agency (CISA), 404, 581–582

D

Dandified YUM (DNF), 304, 316
 Dark Mode, 603
 data at rest, 429
 data bias, 671
 data encryption, as a security best practice, 429
 data governance, 674
 data in transit, 429
 data in use, 429
 data loss prevention (DLP), 387–388
 data protection officer (DPO), 735
 data restoration, 780
 data retention, requirements for, 736–737
 datasets, artificial intelligence (AI), 667–669
 data-usage limit notification, 612
 date code, 209
 dedicated video cards, requirements for, 19
 deep learning, 664
 Default Apps section (Apps Settings screen), 140
 default credentials, changing, 514
 default passwords, 431
 default passwords, changing, 505–506
 default usernames, changing, 437–438, 505–506
 definitions database, 406
 defragmenting, 302
 degaussing, 443–444
 degraded browser performance, 588–589
 degraded performance, 543–548, 612
 Delete option (Advanced tab of System applet), 108
 demilitarized zone (DMZ), 516–517
 demineralized water, 712
 denatured isopropyl alcohol, 712
 denial-of-service (DoS) attacks, 418–420
 deployment, 207–209
 Deployment Image Servicing and Management (DISM), 21
 desktop, 38, 42–46, 582, 659
 destruction and disposal methods, 440–445
 Details tab (Task Manager), 73–75
 detection and response, 407–408
 Dev Drive Protection, 480
 developer mode, 610
 device drivers, rolling back, 576–577
 Device Manager applet, 95, 119–120
 Device Manager snap-in, 81–83
 Device Performance & Health screen (Settings app), 135, 137
 Device Security screen (Settings app), 134
 Devices and Printers applet, 96, 117, 118
 Devices section (Bluetooth & Devices Settings screen), 146, 147

- df command, 320
- DHCP reservation, 518
- Diagnostic & Feedback section (Privacy Settings screen), 141
- Dial-Up section (Network & Internet section), 149
- dictionary attacks, 423
- differential backups, 779
- digital certificates, 586
- digital rights management (DRM), 731
- dir command, 218, 633
- directories, 57, 328–344
- directory services, 388–389
- disabled state, 82
- disappearing files, 584
- Disk Cleanup, 93–95
- disk commands, 230–232
- Disk Defragmenter, accessing, 166
- disk maintenance, 277
- Disk Management, 80–81, 158–166
- Disk tab (Resource Monitor), 92
- Disk Utility, 277
- diskpart command, 231
- disks, 163, 543
- disk-to-cloud backups, 778
- disk-to-disk backups, 778
- disk-to-tape backups, 777
- display devices, 720
- Display section (System section), 144
- displays, 18, 47, 103, 177, 327
- Displays tab (System Settings screen), 280–281
- disposal procedures, 717–722
- distributed denial-of-service (DDoS), 419–420
- DMG filenames, 274
- dnf command, 320
- Dock, 293
- documentation
 - about, 749
 - backup and recovery
 - about, 773
 - backup media, 777–778
 - backup methods, 779
 - backup testing, 780–781
 - critical applications, 777
 - data restoration, 780
 - file-level backups, 773–776
 - image-based backups, 776–777
- change management best practices
 - about, 765
 - change board, 772
 - change type and schedule, 769–770
 - documented business
 - processes, 765–766
 - end-user acceptance, 772–773
 - purpose of change, 767–768
 - request forms, 767
 - risk analysis, 770–772
 - scope of change, 768–769
- exam essentials, 794–795
- for incidents, 727–728
- performance-based question, 800
- performance-based question
 - answers, 835–836
- review question answers, 821–824
- review questions, 796–799
- support and
 - about, 749–750
 - asset management, 755–758
 - common documentation, 758–763
 - knowledge base/articles, 763–764
 - ticketing systems, 750–755
- documented business processes, 765–766
- domain controller, 243
- Domain Information Groper tool, 346–347
- Domain Name System (DNS)
 - about, 256, 259
 - DNS poisoning, 421, 587–588
 - DNS spoofing, 421
 - servers, 256
- domain names, resolving, 346–347
- Domain network profile, 255
- domains, 36, 105, 237–240, 390–391
- door locks, 369–370
- double-click, 43
- download and installation, of web
 - browsers, 489–491

downloadable applications,
 installing, 273–275
 downloadable media, installation and, 21
 drills, for physical destruction, 443
 drive navigation, 219–220
 drivers, 7, 82, 298, 305
 drives, 164, 442
 drop attack, 439
 du command, 343–344
 dumpster diving, 418, 444–445
 dynamic disks, 163
 Dynamic Host Configuration Protocol
 (DHCP), 256, 518–519
 Dynamic Lighting, 45, 137
 dynamic link library (DLL), 549

E

Ease of Access Center applet, 96, 125, 126
 editions
 upgrading, 209
 Windows, 33–35
 educating end users, 597
 educational impact, hallucinations and, 672
 8.3 naming convention, 160
 802.1X tab (System Settings screen), 283
 802.11 standards, 510
 electromagnet, for physical
 destruction, 443–444
 electromagnetic interference (EMI), 694
 electronic health records (EHRs), 736
 electrostatic discharge (ESD), 688–694
 email, as multifactor authentication
 (MFA), 383
 Email & Accounts section (Accounts
 screen), 152
 Email security gateway systems, 408–409
 emergent bias, 671
 employees, 674, 703
 employer responsibilities, for workplace
 safety, 703
 Encrypting File System (EFS), 487–488
 encryption, 473, 507–508, 522
 end users

educating, 597
 end-user acceptance, 772–773
 end-user education, 410–411
 as a security best practice, 432–434
 end-of-life (EOL), 13
 Endpoint Detection and Response (EDR),
 407, 523–524
 end-user license agreement (EULA), 184,
 202–203, 731
 Entra ID, 31–32
 environment variables, 626–627
 Environment Variables button (Advanced tab
 of System applet), 106–107
 environmental concerns
 about, 444, 684, 737
 exam essentials, 737–738
 exercises, 714
 performance-based question answers, 834
 performance-based questions, 743
 policies, licensing, and privacy, 722–737
 real world scenarios, 720–721, 733–734
 review question answers, 819–821
 review questions, 739–742
 environmental controls, 706–722
 Environmental Protection Agency (EPA), 718
 EP cartridge, 688
 equipment, securing, 433
 equipment locks, 370–372
 erosion of trust, hallucinations and, 672
 /etc/fstab configuration file, 322–323
 Ethernet or Wi-Fi section (Network &
 Internet section), 149
 ethical framework, 674
 evaluation bias, 671
 Event Viewer, 79–80, 555
 evil twin, 420
 evolution, artificial intelligence (AI), 661–662
 exam essentials
 administration (Windows), 260–261
 artificial intelligence (AI), 675–676
 documentation, 794–795
 environmental concerns, 737–738
 Linux, 351–352
 macOS, 351–352
 operating system security, 532–533

- operating systems, 62–63
- professionalism, 794–795
- remote access, 675–676
- safety, 737–738
- scripting, 675–676
- security, 445–447
- troubleshooting, 615
- Windows configuration, 167
- exceptions, for firewalls, 485
- Exclusions, 482
- executable machine code, 623
- Executive kernel, 80
- executive subsystems, 550
- exercises
 - administration (Windows), 219–220, 232
 - artificial intelligence (AI), 669–670
 - environmental concerns, 714
 - Linux, 341
 - macOS, 274–275, 276
 - operating system, 46, 47, 49
 - operating system security, 466, 473–474, 490, 521–522
 - professionalism, 787, 791
 - safety, 702
 - scripting, 630–631, 634–635
 - security, 407, 414–415
 - troubleshooting, 556–557
 - Windows configuration, 77, 86, 98, 103
- Experiments menu, 504
- explainability, 674
- exploits and vulnerabilities, 426–428
- explorer.exe, 551
- Extended Detection and Response (XDR), 408
- Extended File Allocation Table (exFAT), 159
- Extended File System (XFS), 159
- extended support stage, in operating system
 - life cycle, 15
- Extensible Firmware Interface (EFI), 324–325
- extensions, 58, 298, 491–492
- external data sources, 667
- external drives, booting from, 215
- Eye Control section (Accessibility Settings screen), 155

F

- F1 key, 55
- Face ID, 521
- face lock, 521
- FaceTime, 292–293
- Facial Recognition Technology (FRT), 378
- factory-reset, 437–438
- failed login attempts, 436–437, 522–523
- fake security warnings, 613–614
- false alerts and hoaxes, 581–582
- Family Options screen (Settings app), 135
- Family section (Accounts screen), 152
- Fast Startup, 114
- FAT16, 160
- FAT32 filesystem, 160
- features
 - macOS, 289–295
 - updating, 209–212
 - web browsers, 502–505
 - Windows, 35–42
- fences, 372–373
- field replaceable units (FRUs), 686
- File Allocation Table (FAT), 158, 160, 466
- File Explorer, 58–60
- File Explorer Options applet, 96, 100–103
- file management. *See also* navigation and file management
 - about, 57
 - capabilities of File Explorer, 58
 - files and folders, 57–58
 - Linux
 - about, 334
 - copying files, 334
 - editing files, 338
 - locating files and content, 338–341
 - moving files, 335–337
 - viewing files, 337
 - navigating File Explorer, 59–61
- file transfer software, 659
- File Vault, 277–278
- file-based backups, 775–776
- fileless, 398–399
- file-level backups, 773–776

- filenames, 58, 97
 - files and folders
 - about, 57–58
 - attributes, 471–474
 - changing file attributes, 478–479
 - collapsing, 59
 - copying, 470–471
 - erasing, 52
 - expanding, 59
 - finding, 59
 - moving, 470–471
 - opening, 59
 - restoring deleted, 51
 - selecting, 59
 - shared, 474–477
 - showing hidden, 103
 - system, 477–488
 - filesystems
 - about, 80
 - Apple (APFS), 159–160
 - information, in order of volatility, 725
 - Linux, 321–324
 - Windows, 160–161
 - find command, 338–341
 - Find My Device section (Privacy Settings screen), 141
 - Find My iPhone app, 525–527
 - Finder app, 275, 293, 295
 - fingerprint locks, 520
 - fingerprint scanners, 377
 - fire safety, 700
 - fire-proof container, 301
 - fire-rated, 301
 - Firewall & Network Protection screen (Settings app), 133, 136
 - firewalls, 121, 252–255, 482–483, 524–525
 - firmware, 305, 514
 - first-in, first-out (FIFO), 778
 - 5 GHz, 509–510
 - fixed tangible assets, 755
 - flash drives, booting from, 215
 - flashlight, 699
 - flexibility, 788–789
 - floating-point, 625
 - Focus section (System section), 144
 - folder redirection, 395
 - folders. *See* files and folders
 - follow-up, 754–755
 - Fonts section
 - Display Settings window, 45
 - Settings app, 138
 - For Developers section (System section), 144–145
 - Force Quit, 278
 - format command, 161, 230–231
 - foundation models, artificial intelligence (AI), 665–667
 - Fourth Extended File System, 159
 - freemium, 272
 - freeware, 731
 - frequencies, SOHO networks (wireless), 509–510
 - frozen system, 601
 - fsck command, 323–324
 - full backups, 779
 - full format, 161
 - full-device encryption, 522
 - fully qualified domain name (FQDN), 226
-
- G**
- G Suite, 29
 - gaffing tape, 702
 - Game Bar section (Gaming Settings screen), 149, 150
 - Game Mode section (Gaming Settings screen), 149
 - Gaming Settings screen (Settings app), 149–150
 - gateway, 256–257, 259
 - General Availability Channel, 211
 - General Data Protection Regulation (GDPR), 735
 - General section (Privacy Settings screen), 141
 - General tab
 - Folder Options applet, 101, 102
 - System Configuration tool, 88, 89

gestures, 293, 294
 globally unique identifier (GUID), 243, 461
 Google Chat, 31
 Google Chrome, 492
 Google Docs, 30
 Google Drive, 528
 Google Meet, 30
 Google Neural Machine Translation (GNMT), 662
 Google Sheets, 30
 Google Slides, 30–31
 Google Workspace, 29–30
`gprestart` command, 228
 GPT disks, 163
`gpupdate` command, 227–228
 GRand Unified Bootloader (GRUB), 324–325
 grandfather, father, son (GFS), 778
 graphical user interface (GUI), 8
 graphics card, system requirements for, 177
 graphics performance, 543
 Group Policy, 22, 38, 392–393
 Group Policy Editor snap-in, 86–88
 Group Policy Management Console (GPMC), 393
 Group Policy Objects (GPOs), 228, 239, 255, 389, 391, 392
 groups, 314–315, 457–460
 guest network isolation, SOHO networks (wireless), 506
 Guest user account, 438, 459

H

hallucinations, 671–672
 hammers, for physical destruction, 443
 hard disks, 166
 hard drives, 176, 441–442
 hardening techniques

- about, 520
- configuration profiles, 525
- endpoint detection and response (EDR), 523–524
- failed login attempts restrictions, 522–523
- firewalls, 524–525
- full-device encryption, 522
- screen locks, 520–522

 hardware, 15–17, 177
 Hardware Abstraction Layer (HAL), 550
 Hardware Compatibility Lists (HCLs), 16, 177
 Hardware tab

- System applet, 104
- System Settings screen, 283

 hardware tokens, as multifactor authentication (MFA), 383–385
 hazardous waste, 706
 HDDerase, 442
 healthcare risks, hallucinations and, 672
 Hearing Devices section (Accessibility Settings screen), 155
 Help, accessing, 55
`help` command, 234
 hex drivers, 698
 Hibernate state, 113
 Hidden attribute, 471
 high number of ads, 613
 high resource utilization, 611–612
 high-performance computing (HPC), 33
 high-voltage probe, 687
 Hiren's BootCD, 442
 historical bias, 670
 hives, 156
 HKEY_CLASSES_ROOT hive, 156
 HKEY_CURRENT_CONFIG hive, 156
 HKEY_CURRENT_USER hive, 156
 HKEY_LOCAL_MACHINE hive, 156
 HKEY_USERS hive, 156
 home directory, 329
 home folders, 395
 host-based firewalls, 121
`hostname` command, 232–233
 HOSTS file, 347
 hot swap disks, 164
 hot swappable drives, booting from, 215
 hybrid attacks, 423
 hybrid identity, 31–32
 Hypertext Markup Language (HTML), 639–640
 hyphen (-), 331

-
- I**
 - iCloud, 291–293, 526, 528
 - iCloud Drive, 291
 - icons
 - about, 50
 - Microsoft Edge, 53
 - pinning to Start menu, 48
 - Recycle Bin, 50–53
 - standard desktop, 50–53
 - This PC, 53
 - identification (ID) badges, 367
 - identity and access management (IAM), 31, 388
 - ifconfig command, 344–345
 - image deployment, 21, 207
 - image drum, 688
 - image recovery, 277
 - image-based backups, 776–777
 - iMessage, 291–292
 - impact to business, 23
 - impact to device, 23
 - impact to network, 23
 - impact to operation, 23
 - impersonation, 417
 - incident response, 724–729
 - incidents, 705–706, 759–760
 - incineration, for physical destruction, 444
 - incremental backups, 779
 - Indexing Options applet, 96, 123–125
 - Indexing Service, 472
 - information gathering, scripting
 - and, 645–646
 - informational tools
 - /? switch, 234
 - about, 232
 - help command, 234
 - hostname command, 232–233
 - securing, 433
 - whoami command, 233–234
 - winver command, 233
 - inheritance, 475–477
 - initiating updates, scripting and, 646–647
 - Inking & Typing Personalization section (Privacy Settings screen), 141
 - in-person testing, 772
 - in-place restoration, 780
 - InPrivate browsing, 498
 - Insider Program, 211
 - insider threats, 423–424
 - Installed Apps section (Apps Settings screen), 139–140
 - installing
 - applications
 - about, 20–22
 - from App Store, 272–273
 - scripting and, 644–645
 - downloadable applications, 273–275
 - software packages, 761–762
 - Windows
 - about, 176–179
 - boot options, 212–216
 - considerations, 216–217
 - options for, 179–180
 - process for, 180–200
 - recovery partition, 206
 - repair installation, 205–206
 - instant messaging, 31
 - integers, 625
 - integrated circuit chip (ICC), 375
 - integrated development environment (IDE), 638
 - Integrated Scripting Environment (ISE), 633, 638
 - integration, artificial intelligence (AI), 669–670
 - interactive commands, 327
 - interface, Windows
 - about, 42
 - Desktop, 42–46
 - icons, 50–53
 - Start menu, 47–50
 - Taskbar, 46–47
 - intermittent wireless, 604–605
 - internal hard drives, booting from, 215–216
 - internal service-level agreement (SLA), 763

International Organization for
 Standardization (ISO), 713, 760
 Internet connectivity, limited, 612–613
 Internet Control Message Protocol
 (ICMP), 223
 Internet Explorer, 100
 Internet of Things (IoT) considerations, 531
 Internet Options applet, 96, 99–101
 Internet Protocol (IP) addressing scheme, 256
 Internet Protocol (IP) filtering, 514–515
 Internet-based boot, 215
 intrusion detection system (IDS), 724–725
 inventory lists, 756
 iOS, 12
 iOS App Store Package (IPA), 609
 IP addressing, 259
 iPadOS, 12
 ipconfig command, 223, 344–345
 iPhones, setting passcode locks on, 521
 ISO mountables, installation and, 20

J

Java, 625
 Java Runtime Environment (JR), 625
 JavaScript, 639–640
 journaled file system (JFS), 301
 just-in-time access, 386

K

Kerberos, 513
 Kerberos token, 461
 kernel initialization, 213
 key fobs, 374, 384
 Keyboard section (Accessibility Settings
 screen), 155
 keyboards, 688
 Keychain Access, 290
 keyloggers, 397–398
 keys, 375–376
 knowledge base/articles, 763–764

L

languages, scripting
 JavaScript, 639–640
 PowerShell scripts, 631–635
 Python scripts, 638–639
 Shell scripts, 636–637
 Visual Basic scripts, 635–636
 Windows batch scripts, 629–631
 large language models (LLMs), 666
 laser printers, toner disposal, 721
 Last Known Good Configuration option, 157
 leaked personal files/data, 614
 legal and financial, hallucinations and, 672
 /Library folder, 297–298
 licensing, 19, 32, 755–756
 lid, closing the, 113–114
 life cycle, of operating systems, 13–15
 lighting, 378
 Lightweight Directory Access Protocol
 (LDAP), 388
 limitations, artificial intelligence
 (AI), 670–673
 limited Internet connectivity, 612–613
 Linear Tape-Open (LTO), 777
 link, 330
 Linux
 about, 11, 271–272, 306–307, 350–351
 commands, 309–312
 directory listings, 330–334
 directory management, 342–344
 directory navigation, 328–330
 exam essentials, 351–352
 exercises, 341
 file management, 334–341
 groups, 314–315
 maintaining operating system, 315–325
 networking utilities, 344–350
 passwords, 314–315
 performance-based question
 answers, 831–832
 performance-based questions, 357
 process management, 325–328
 review question answers, 808–810

- review questions, 353–356
- syntax of shell commands, 307–309
- user accounts, 312–315
- Linux filesystems, 158–159
- liquid crystal display (LCD), 687
- lithium-ion (Li-ion) batteries, 720
- load average, 327–328
- local library, 297
- Local shares, 474–475
- Local Users and Groups snap-in, 84–85
- location-based services, 603
- Lock Screen
 - Display Settings window, 45
 - Settings app, 137
- Lockdown Mode, 285–286
- logging, 392
- logging off, 433
- Logical Disk Management (LDM), 163
- logical security
 - about, 379
 - access control lists (ACLs), 380–382
 - Active Directory (AD), 389–396
 - authentication factors, 382
 - data loss prevention (DLP), 387–388
 - directory services, 388–389
 - identity and access management (IAM), 388
 - just-in-time access, 386
 - mobile device management (MDM), 387
 - multifactor authentication (MFA), 382–385
 - principle of least privilege, 379
 - privileged access management (PAM), 387
 - single sign-on (SSO), 385–386
 - zero trust, 379–380
- Logical Volume Manager (LVM), 322
- login scripts, 393–394
- logs, Event Viewer, 79
- Long-Term Servicing Channel, 211
- long-term support (LTS), 13, 15
- loop device, 322
- loops, scripting, 627–628
- low memory warnings, 561–563

- low-level format, 441
- ls command, 331, 333

M

- MAC address filtering, 515
- MAC filtering, wireless, 508
- machine learning (ML), 662–665
- macOS
 - about, 271–272, 350–351
 - applications on, 272–276
 - best practices, 299–306
 - components
 - about, 276
 - features, 289–295
 - system settings, 279–288
 - tools, 277–279
 - exam essentials, 351–352
 - exercises, 274–275, 276
 - performance-based question
 - answers, 831–832
 - performance-based questions, 357
 - review question answers, 808–810
 - review questions, 353–356
 - system folders, 295–299
- macOS Keychain Access app, 434
- macro virus, 402
- magnetometers, 378–379
- magnets, as a source of interference, 694
- Magnifier section (Accessibility Settings screen), 154
- Mail applet, 123
- Mail (Microsoft Outlook) applet, 96
- mainstream support stage, in operating
 - system life cycle, 14–15
- maintenance window, 769
- malicious applications, 611
- malicious plug-ins, 588–589
- malware
 - about, 396
 - adware, 404–405
 - best practices for removal of, 589–597
 - botnets, 404

- cryptominers, 398
- fileless, 398–399
- keyloggers, 397–398
- ransomware, 396–397
- rootkits, 398
- spyware, 398
- stalkerware, 399
- Trojan horse, 397
- virus transmission in networks, 403–404
- viruses, 399–403
- worms, 404
- man command, 350
- Managed Detection and Response (MDR), 407–408
- man-in-the-middle (MitM) attack, 422
- manual state, 83
- Master Boot Record (MBR), 162–163, 402
- Master File Table (MFT), 158
- material safety data sheets (MSDSs), 705, 714–717
- Maximize button, in windows, 54
- maximized windows, 56
- md command, 218–219
- media converter, 368
- memory, 18, 561–563, 725
- Memory tab (Resource Monitor), 92, 94
- Menu bar, in windows, 54
- Mercury-Containing and Rechargeable Battery Management Act, 719
- metered connection, 247–248
- Microsoft 365 Apps suite, 26, 30, 32
- Microsoft Active Directory (AD), 22, 31, 388
- Microsoft Defender, 38–40, 593
- Microsoft Disk Operating System (MS-DOS), 217
- Microsoft Download Center, 178
- Microsoft Edge, 37, 100, 200, 201
- Microsoft Edge icon, 53
- Microsoft Endpoint Configuration Manager (MECM), 22, 207, 209, 426, 575
- Microsoft Entra ID, 240–242
- Microsoft Excel, 30
- Microsoft Exchange, 25–28, 32
- Microsoft Management Console (MMC)
 - about, 77–78
 - Certificate Manager snap-in, 83–84
 - Computer Management snap-in, 78–79
 - Device Manager snap-in, 81–83
 - Disk Management, 161–165
 - Disk Management snap-in, 80–81
 - Event Viewer snap-in, 79–80
 - Group Policy Editor snap-in, 86–88
 - Local Users and Groups snap-in, 84–85
 - Performance Monitor snap-in, 85–86
 - Task Scheduler snap-in, 81
- Microsoft OneDrive, 28–29, 32, 197–200
- Microsoft Outlook, 25–28
- Microsoft PowerPoint, 30–31
- Microsoft Store, 33
- Microsoft Teams, 30, 31, 32
- Microsoft Windows Tablet, 11–12
- Microsoft Word, 30
- microwave ovens, as a source of interference, 694
- Minimize button, in windows, 54
- minimized windows, 56
- mirrors, 699
- misinformation, hallucinations and, 672
- Mission Control, 289–290
- mitigating software threats
 - about, 405
 - anti-malware, 406–407
 - antivirus, 405–406
 - detection and response, 407–408
 - Email security gateway systems, 408–409
 - end-user education, 410–411
 - recovery console, 409–410
 - reinstalling OS, 413
 - software firewalls, 411–413
- mitigation strategies, 672–673
- mkdir command, 342
- mobile device management (MDM), 13, 273, 387, 428, 525, 530, 610, 614
- mobile devices
 - disposal of, 721–722
 - security

- about, 520
- authenticator applications, 529
- hardening techniques, 520–525
- Internet of Things (IoT)
 - considerations, 531
- multifactor authentication, 529
- patch management, 528–529
- policies and procedures, 530–531
- remote backup, 528
- remote wipes and locator
 - applications, 525–527
- Mobile Devices section (Bluetooth & Devices Settings screen), 147
- mobile digital key, 376
- mobile operating systems
 - defined, 11
- troubleshooting
 - about, 597
 - application problems, 597–600
 - autorotate issues, 607–608
 - battery life, 603–604
 - connectivity issues, 604–607
 - OS Fails to Update, 602
 - performance issues, 600–601
 - random reboots, 601–602
- mobile security issues, troubleshooting
 - about, 608
 - developer mode, 610
 - fake security warnings, 613–614
 - high number of ads, 613
 - high resource utilization, 611–612
 - leaked personal files/data, 614
 - limited Internet connectivity, 612–613
 - malicious applications, 611
 - root access/jailbreak, 610
 - unexpected application behavior, 614
 - unofficial sources of applications, 609
- monitors, 686–687, 720–721
- motion sensors, 368–369
- mount command, 321–322
- mounting partitions, 163
- mouse, 688
- Mouse Pointer and Touch section (Accessibility Settings screen), 154

- Mouse section
 - Accessibility Settings screen, 155
 - Bluetooth & Devices Settings screen, 147
- msinfo32 command-line options, 92
- msinfo32.exe command, 177–178
- multibooting, 216
- multifactor authentication (MFA), 382–385, 512, 529
- multimeter, 699
- multipartite virus, 402
- Multitasking section (System section), 144
- multithreading, 9
- multiuser license, 732
- mutual nondisclosure agreement (MNDA), 729
- mv command, 335–337

N

- N version, 34
- Narrator section (Accessibility Settings screen), 154
- National Institute of Standards and Technology (NIST), 430
- natural language processing (NLP), 665
- navigation and file management
 - about, 218
 - cd command, 218–219
 - copy command, 220–221
 - dir command, 218
 - drive navigation, 219–220
 - File Explorer, 59–60
 - md command, 218–219
 - rd command, 218–219
 - robocopy command, 222
 - xcopy command, 221–222
- Nearby Sharing section (System section), 144
- near-field communication (NFC), 605
- Nessus vulnerability scanner, 304
- net command, 227
- netstat command, 226
- network access token, 461

- Network and Sharing Center applet, 96, 119, 120
 - network devices, as a source of interference, 694
 - network drives, remapping, 644
 - Network File System (NFT), 159
 - Network & Internet section (Settings app), 148–149
 - network processes, in order of volatility, 725
 - network shares, 250–251
 - Network tab
 - Resource Monitor, 92–93
 - System Settings screen, 282
 - Network Time Protocol (NTP) server, 128
 - network traffic, 612
 - networks and networking
 - booting from, 214–215
 - connectivity issues, 581
 - connectivity tools
 - about, 223, 347–350
 - ipconfig command, 223
 - net command, 227
 - netstat command, 226
 - nslookup command, 226–227
 - pathping command, 224–225
 - ping command, 223–224
 - tracert command, 224
 - defined, 9
 - establishing connections, 244–249
 - installing, 21–22
 - Linux
 - getting help, 350
 - HOSTS file, 347
 - managing network interfaces, 344–345
 - network connectivity, 347–350
 - resolvers, 345–346
 - resolving domain names, 346–347
 - managing interfaces, 344–345
 - models
 - about, 236
 - domains, 237–240
 - Microsoft Entra ID, 240–242
 - workgroups, 236–237
 - performance, 543
 - virus transmission in, 403–404
 - in Windows
 - about, 236
 - accessing resources, 249–252
 - client network configuration, 256–259
 - establishing connections, 244–249
 - firewall settings, 252–255
 - models for, 236–242
 - user authentication, 243–244
 - neural networks, 663–664
 - New Technology Filesystem (NTFS), 158, 160, 466–474
 - new-user setup checklist, 762
 - nickel-cadmium (NiCd) batteries, 719
 - nickel-metal hydride (NiMH) batteries, 720
 - “No OS Found,” 551–552
 - noncompliant systems, as a vulnerability, 426
 - nondisclosure agreement (NDA), 729
 - normative ethics, 791
 - Notepad, 55
 - Notifications section (System section), 144
 - nslookup command, 346
 - nslookup command, 226–227
 - ntbtlog.txt, 549
 - ntoskrnl.exe, 549
-
- O**
- objects, 50–51, 60, 391–392
 - Occupational Safety and Health Act (OSHA), 702–706
 - offboarding procedure, 530, 762
 - Offline Maps section (Apps Settings screen), 141
 - off-site storage, 302
 - onboarding procedure, 530, 762
 - One-Time Password/Passcode (OTP), 384
 - on-path attacks, 422
 - on-premises Email security gateways, 408
 - on-premises productivity, 24–25
 - on-site storage, 301
 - Open security, 507

- open source, 8, 731
- operating system (OS)
 - about, 6–8, 61–62
 - applications
 - about, 17
 - cloud-based productivity tools, 23–32
 - considerations for, 23
 - installation, 20–22
 - security considerations, 22–23
 - system requirements for, 18–20
 - categories, 10–11
 - common, 11–13
 - exam essentials, 62–63
 - exam preparation, 61
 - exercises, 46, 47, 49
 - interacting with
 - about, 72
 - Control Panel, 95–127
 - Disk Cleanup, 93–95
 - Microsoft Management Console (MMC), 77–88
 - Resource Monitor, 92–93
 - System Configuration tool, 88–90
 - System Information tool, 90–92
 - Task Manager, 72–77
 - Windows Settings app, 127–155
 - life cycle, 13–15, 427
 - maintaining, 315–325
 - management tools
 - about, 227
 - gpresult command, 228
 - gpupdate command, 227–228
 - sfc command, 228–229
 - shutdown command, 229–230
 - minimum requirements, 15–17
 - notifications, 582–583
 - performance-based question answers, 826
 - performance-based questions, 67
 - reimaging, 574–575
 - reinstalling, 413, 574–575
 - review question answers, 802–803
 - review questions, 64–67
 - stateless, 12–13
 - terms and concepts, 8–10
 - update failures, 584–585
 - updating, 528–529
 - Windows
 - about, 32
 - editions, 33–35
 - elements of windows, 53–55
 - features, 35–42
 - file management, 57–61
 - interface, 42–53
 - states of windows, 56
- operating system security
 - about, 532
 - exam essentials, 532–533
 - exercises, 466, 473–474, 490, 521–522
 - mobile devices, 520–531
 - performance-based question
 - answers, 832–833
 - performance-based questions, 538
 - real world scenarios, 500, 527
 - review question answers, 812–814
 - review questions, 534–537
 - settings
 - about, 457
 - New Technology File System (NTFS), 466–474
 - share permissions, 466–474
 - shared files and folders, 474–477
 - system files and folders, 477–488
 - user authentication, 460–466
 - users and groups, 457–460
 - small office, home office (SOHO)
 - networks (wired), 513–519
 - small office, home office (SOHO)
 - networks (wireless), 505–513
 - web browser, 488–505
- operational inefficiencies, hallucinations
 - and, 672
- optimizing hard disks, 166
- Optional Features section (System
 - section), 145
- options, 307
- order of volatility, 725–726

organizational units, 391
 original equipment manufacturer (OEM), 213
 OS Fails to Update, in mobile OS, 602
 Other Users screen (Accounts screen), 152
 out-of-box drivers, 576–577
 out-of-box experience pass, 191
 overheating, 603–604
 oversight, 674
 overwriting drives, 442, 780
 ownership, changing, 332–334

P

page/paging file, 562
 palm print scanners, 377–378
 pan-tilt-zoom (PTZ), 367
 parentheses, filenames in, 97
 partitioning, 158
 partitions, 163, 164, 215–216
 passive detection, 724
 passive infrared (PIR), 368–369
 Passkeys section (Accounts screen), 152
 password attacks, 422–423
 Password Complexity policy, 430
 passwords, 314–315, 432, 434, 439, 444–445
 Passwords app, 290
 patch management, 304, 528–529
 patching, 427, 491
 path, 57
 pathping command, 224–225
 pattern lock, 521
 Payment Card Industry Data Security
 Standard (PCI DSS), 735
 PC toolkit, 697
 peer review, 770
 Pen & Window Ink section (Bluetooth &
 Devices Settings screen), 148
 penetration testing, 725
 pentesting, 725
 performance, 543–548, 600–601
 Performance Monitor snap-in, 85–86
 Performance option (Advanced tab of System
 applet), 106
 Performance tab (Task Manager), 76
 performance-based question answers
 administration (Windows), 827–830
 artificial intelligence (AI), 834
 configuration (Windows), 826–827
 documentation, 835–836
 environmental concerns, 834
 Linux, 831–832
 macOS, 831–832
 operating system (OS), 826
 operating system security,
 832–833
 professionalism, 835–836
 remote access, 834
 safety, 834
 scripting, 834
 security, 832
 troubleshooting, 833
 performance-based questions
 administration (Windows), 266
 artificial intelligence (AI), 680
 documentation, 800
 environmental concerns, 743
 Linux, 357
 macOS, 357
 operating system security, 538
 operating systems, 67
 professionalism, 800
 remote access, 680
 safety, 743
 scripting, 680
 security, 452
 troubleshooting, 619
 Windows configuration, 171
 period (.), 334
 permissions
 about, 392
 changes in, 584
 effective, 469–470
 inheritance, 475–477
 Linux, 331–334
 propagation, 475–477
 share, 466–474
 perpetual license, 732

- personal government-issued information, 735
- personal identification number (PIN),
 - 382, 463
- Personalization Settings screen (Settings app), 136–139
- personally identifiable information (PII), 734
- phage virus, 402–403
- pharming, 587–588
- phishing, 414, 415
- Physical Address Extension (PAE), 10
- physical destruction, 443–445
- physical documents, securing, 444–445
- physical environment, managing, 706–713
- physical media, installation and, 20
- physical security
 - about, 366
 - access control vestibule, 366
 - alarm system, 369
 - badge reader, 367
 - bollards, 373–374
 - door locks, 369–370
 - equipment locks, 370–372
 - fences, 372–373
 - motion sensors, 368–369
 - security guards, 372
 - SOHO networks (wired), 519
 - for staff
 - about, 374
 - biometrics, 376–378
 - key fobs, 374
 - keys, 375–376
 - lighting, 378
 - magnetometers, 378–379
 - mobile digital key, 376
 - RFID badges, 375
 - smartcards, 375
 - video surveillance, 367–368
- picture password, 464
- PIN lock, 521
- ping command, 223–224, 348
- pinning, 41, 48
- pipe (|), 309
- PKG filenames, 274
- plagiarism, 675
- pliers, 699
- policies, 389, 530–531, 673–674
- polymorphic virus, 403
- pop-up blockers, 495–496
- pop-ups, 585–586
- port forwarding, 517–518
- port mapping, 517–518
- ports, disabling, 515
- post-installation setup, 196–200
- potentially unwanted programs (PUPs), 404–405
- power, 113, 707–710
- Power Options applet, 96, 111–115
- Power Saving Mode, enabling, 603
- Power section (System section), 144
- power strips, 707
- power supply, 685–686
- Power Users group, 459
- Power-On Self-Test (POST), 213, 548
- PowerShell, 626
- PowerShell scripts, 631–635
- Preboot Execution Environment (PXE),
 - 208, 212
- preemptive multitasking, 9
- prerequisites, requirements for, 18
- presentation tools, 30–31
- primary mouse click, 43
- principle of least privilege, 379, 434
- printer shares, 251
- printers, 687–688
- Printers & Scanners section (Bluetooth & Devices Settings screen), 146
- Printers & Scanners tab (System Settings screen), 284–285
- privacy, 667–669, 791–793
- Privacy & Security tab (System Settings screen), 285
- Privacy Settings screen (Settings app), 141–143
- Private network profile, 255
- private-browsing mode, 498
- privileged access management (PAM), 387

Privileged Identity Management (PIM), 387
 proactive detection, 725
 problem state, 82
 process management, in Linux, 325–328
 Process tab (Task Manager), 73
 processors, system requirements for, 176
 procurement life cycle, 757–758
 product key, 732
 professionalism
 about, 749, 781, 793–794
 appearance, 781–782
 appropriate behavior, 787–793
 customer communication, 782–786
 exam essentials, 794–795
 exercises, 787, 791
 performance-based question, 800
 performance-based question
 answers, 835–836
 real world scenarios, 784, 786, 792
 review question answers, 821–824
 review questions, 796–799
 profiles, 255, 564–565
 program variables, 627
 program window, 53
 programs, starting from Run window, 49
 Programs and Features applet, 95
 Programs and Features applet, 116–117
 progress notes, 754
 prohibited content/activity, 723–724
 Projecting To This PC section (System
 section), 145
 protected health information (PHI), 736
 Protection History screen (Settings app),
 135, 138
 Proxies tab (System Settings screen), 283
 Proxy section (Network & Internet
 section), 149
 proxy servers, 248–249, 500–501
 Public network profile, 255
 punctuality, 787
 purpose of change, 767–768
 push installations, 22
 Python scripts, 638–639

Q

QR code phishing, 416
 quality management system (QMS), 760–761
 quarantining infected systems, 591
 question mark (?), 59, 221
 Quick Assist, for remote access, 657–659
 quick format, 161
 quishing, 416

R

radio frequency identification
 (RFID), 374–375
 radio power levels, SOHO networks
 (wireless), 511
 radio-frequency interference (RFI), 694
 rainbow tables, 423
 random access memory (RAM), 10, 176,
 543, 561–563
 random reboots, mobile OS, 601–602
 ransomware, 396–397
 Rapid Security Response (RSR), 305–306
 raw disk blocks, in order of volatility, 725
 rd command, 218–219
 Read-only attribute, 471
 real world scenarios
 environmental concerns,
 720–721, 733–734
 operating system security, 500, 527
 professionalism, 784, 786, 792
 safety, 690, 695, 703
 security, 445
 troubleshooting, 561
 Windows configuration, 157–158
 real-time clock (RTC), 565
 real-time data, 667
 real-time protection, 480
 rebooting, as a troubleshooting step, 567
 Recommended section, of Start menu, 50
 recovery console, 409–410
 recovery partition, 206

- Recovery screen (System section), 145
- recovery time objective (RTO), 780
- Recycle Bin icon, 50–53
- recycling/repurposing best practices, 441–442
- Red Hat Packager Manager (RPM)-based, 304
- reduced functionality mode (RFM), 181
- Redundant Array of Independent (or Inexpensive) Disks (RAID), 164–165
- reflective DoS, 418
- Regional and Language Options, 99
- Registry Editor, 156–158
- regulations, 730–731
- regulatory concerns, 444
- reimaging OS, 574–575
- reinforcement learning, 663
- reinstalling OS, 413, 574–575
- release candidate stage, in operating system life cycle, 14
- release to market (RTM), 14
- Reliability Monitor, 555–557
- remapping network drives, scripting and, 644
- remediating, 592–595, 727
- remote access
 - about, 623, 649–650, 675
 - exam essentials, 675–676
 - performance-based question answers, 834
 - performance-based questions, 680
 - Remote Desktop Protocol (RDP), 650–651
 - Remote Monitoring and Management (RMM), 655–656
 - review question answers, 817–819
 - review questions, 677–679
 - Secure Shell (SSH), 654–655
 - security considerations, 659–660
 - Simple Protocol for Independent Computing Environments (SPICE), 656
 - tools for, 657–659
 - Virtual Network Computing (VNC), 653–654
 - Virtual Private Network (VPN), 651–653
 - Windows Remote Management (WinRM), 657
- Remote Assistance tool, 110
- Remote Authentication Dial-In User Service (RADIUS), 513
- Remote Backup, 528
- Remote Desktop Protocol (RDP), 36, 650–651
- Remote Desktop section (System section), 145
- Remote Monitoring and Management (RMM), 655–656
- remote network installation, 212
- remote office, 302
- Remote Server Administration Tools (RSAT), 641
- Remote tab (System applet), 110
- remote wipes and locator applications, 525–527
- renamed system files, 584
- repair installation, 205–206
- repair tools, 696–700
- replication, 392
- request forms, 767
- reserved word, 626
- Resilient File System (ReFS), 160
- resolvers, Linux, 345–346
- resource allocation, 326–327
- Resource Monitor, 92–93, 545–547
- resources, 249–252, 611–612
- respect, 790–791
- responsible staff members, 766
- Restore button, in windows, 54
- restore points, 109, 572–574, 596
- restored windows, 56
- retina scanners, 377
- retrovirus, 403
- review question answers
 - administration (Windows), 805–808
 - artificial intelligence (AI), 817–819
 - configuration (Windows), 803–805
 - documentation, 821–824
 - environmental concerns, 819–821
 - Linux, 808–810
 - macOS, 808–810
 - operating system (OS), 802–803

- operating system security, 812–814
- professionalism, 821–824
- remote access, 817–819
- safety, 819–821
- scripting, 817–819
- security, 810–812
- troubleshooting, 815–817
- review questions
 - administration (Windows), 262–265
 - artificial intelligence (AI), 677–679
 - documentation, 796–799
 - environmental concerns, 739–742
 - Linux, 353–356
 - macOS, 353–356
 - operating system security, 534–537
 - operating systems, 64–67
 - professionalism, 796–799
 - remote access, 677–679
 - safety, 739–742
 - scripting, 677–679
 - security, 448–451
 - troubleshooting, 616–618
 - Windows configuration, 168–171
- RFID badges, 375
- risk analysis, 770–772
- risk levels, 770–771
- `rmdir` command, 342
- `robocopy` command, 222
- rogue antivirus, 583–584
- rollback plan, 765–766
- rolling back device drivers, 576–577
- rolling upgrade, 206
- root access/jailbreak, 610
- root folder, 58
- root user, 312–313
- rootkits, 398
- Run window, 49
- exam essentials, 737–738
- exercises, 702
- identifying potential hazards, 684–695
- implementing policies and procedures, 702–706
- performance-based question answers, 834
- performance-based questions, 743
- policies, licensing, and privacy, 722–737
- real world scenarios, 690, 695, 703
- review question answers, 819–821
- review questions, 739–742
- safety goggles, 700
- safety plans, 704
- sags, 707
- sandbox testing, 766
- sanitizing hard drives, 441–442
- Sarbanes-Oxley Act (SOX) of 2002, 730
- scans, scheduling, 595–596
- scheduled backups, 299–302
- scheduled disk maintenance, 302
- schema partition, 389
- scope of change, 768–769
- screen locks, 433, 520–522
- Screen Saver (Display Settings window), 45
- screen savers, 46, 439
- screen sharing, 278, 659
- screened subnet, 516–517
- screwdrivers, 698
- scripting
 - about, 623–625, 647–648, 675
 - branch logic, 629
 - changing system settings, 648–649
 - comment syntax, 627
 - environment variables, 626–627
 - exam essentials, 675–676
 - exercises, 630–631, 634–635
 - languages, 631–640
 - loops, 627–628
 - performance-based question answers, 834
 - performance-based questions, 680
 - review question answers, 817–819
 - review questions, 677–679
 - security, 648
 - system crashes, 649

S

- safety
 - about, 684, 737
 - creating safe workplaces, 695–706
 - documentation for, 714–717

- use cases, 640–647
- variables, 625–626
- Search Permissions section (Privacy Settings screen), 142
- Searching Windows section (Privacy Settings screen), 142
- secondary mouse click, 43
- Secure Boot, 179
- secure data transfers, 494–495
- Secure DNS, 501–502
- Secure Hash Algorithm (SHA), 490
- secure management access, SOHO networks (wired), 519
- Secure Shell (SSD), 654–655
- security
 - about, 365–366, 445
 - of applications, 22–23
 - best practices, 428–440
 - common threats, 418–426
 - datasets and, 667–669
 - destruction and disposal
 - methods, 440–445
 - exam essentials, 445–447
 - exercises, 407, 414–415
 - exploits and vulnerabilities, 426, 427, 428
 - hallucinations and, 672
 - logical, 379–396
 - malware, 396–405
 - mitigating software threats, 405–413
 - performance-based question, 452
 - performance-based question answers, 832
 - physical, 366–374
 - physical, for staff, 374–379
 - real world scenarios, 445
 - remote access, 659–660
 - review question answers, 810–812
 - review questions, 448–451
 - scripting, 648
 - social engineering attacks, threats, and vulnerabilities, 414–418
 - troubleshooting, 580–589
 - updating, 209–212
- Security Account Manager (SAM), 243, 460
- security groups, 395–396
- security guards, 372
- Security Identifier (SID), 234
- security key, 463
- self-grounding, 692–693
- self-paced testing, 772
- semi-structured data, 667
- sensitive information, managing, 734–737
- server locks, 371–372
- server operating systems, 10
- servers, water and, 695
- service pack stage, in operating system life cycle, 14
- service-level agreement (SLA), 762–763
- Services tab (System Configuration tool), 76, 88, 89, 90
- servicing channels, 211–212
- Session Manager, 550
- Session Manager Subsystem, 214
- setting(s)
 - for add-ons, 502–504
 - encryption, 507–508
 - passcode locks on Android phones, 522
 - passcode locks on iPhones, 521
 - strong passwords, 429–430
 - time restrictions, 435
 - web browsers, 495–502
- Settings app
 - about, 98, 127–128
 - Accessibility Settings screen, 153–155
 - Account Protection screen, 132–133, 135
 - Accounts section, 151–152
 - App & Browser Control screen, 133–134
 - Apps Settings screen, 139–141
 - Background section, 136
 - Bluetooth & Devices Settings screen, 145–148
 - Colors section, 136
 - configuring network addressing, 257
 - Device Performance & Health screen, 135, 137
 - Device Security screen, 134
 - Dynamic Lighting, 137
 - Family Options screen, 135
 - Firewall & Network Protection screen, 133, 136
 - Fonts section, 138

- Gaming Settings screen, 149–150
- Lock Screen, 137
- Network & Internet section, 148–149
- Personalization Settings screen, 136–139
- Privacy Settings screen, 141–143
- Protection History screen, 135, 138
- Start section, 138
- System section, 143–145
- Taskbar section, 138
- Text Input section, 138
- Themes section, 137
- Time & Language screen, 128–129
- Virus & Threat Protection screen, 132, 134
- Windows Security screen, 131–136
- Windows Update, 129–131
- setx command, 648–649
- sfc command, 228–229
- share permissions, 466–474
- shared resources, 252
- shareware, 731
- shell, 8, 278
- shell commands, 307–309
- Shell scripts, 636–637
- Short Message Service (SMS), 383, 416
- shortcuts, creating, 276
- shoulder surfing, 417
- shredders, 443, 444–445
- shrinking partitions, 164
- Shut Down command, 49–50
- shutdown command, 229–230
- shutdowns, frequent, 558–559
- Shutdown/Turn Off state, 113
- side-by-side upgrading, in Windows 10/11, 206–207
- sideloading, 609
- signature, 403
- Sign-in Options section (Accounts screen), 152
- signing in, 197–200, 201
- Simple Protocol for Independent Computing Environments (SPICE), 656
- single dot (.), 329
- single sign-on (SSO), 244, 385–386, 462–463
- single user license, 732
- single-factor authentication, 512
- site-to-site VPN, 651–652
- 6 GHz, 510
- 64-bit operating system, 9
- slash (/), 330
- Sleep/Suspend state, 113
- slow-loading profiles, 564–565
- small office, home office (SOHO) networks, 575
- small office, home office (SOHO) networks (wired), security of, 513–519
- small office, home office (SOHO) networks (wireless), security of, 505–513
- smartcards, 375
- smishing, 416
- Snap Assist, 41–42
- snapshots, 277
- social engineering attacks, threats, and vulnerabilities, 414–418
- software
 - for asset management systems, 755–756
 - firewalls, 411–413
 - installing packages, 761–762
 - managing licenses for, 731–734
 - mitigating threats, 405–413
 - traditional, 25
- Software Asset Management (SAM), 733
- software tokens, as multifactor authentication (MFA), 383–385
- SolarWinds attack, 425–426
- solid-state drives (SSDs), 160, 215, 442
- something you are, 382
- something you do, 382
- something you have, 382
- something you know, 382
- somewhere you are, 382
- Sound applet, 96, 117, 118
- Sound section (System section), 144
- source code, 8
- spear phishing, 416
- specialize pass, 188

- Speech section
 - Accessibility Settings screen, 155
 - Privacy Settings screen, 141
- speed, frequency and, 510
- spikes, 707
- splash screens, 730–731
- spoofing attacks, 421–422
- Spotlight, 290–291
- spreadsheets, 30
- spyware, 398
- SSID, changing, 506
- SSID broadcast, disabling, 508
- stakeholder engagement, 674
- stalkerware, 399
- standard desktop icons, 50–53
- standard format, 441
- standard operating procedures (SOPs), 760–761
- standard privileges, commands with, 235
- Standard user account, 460, 465–466
- Standby state, 113
- Start menu, 47–50
- Start section
 - Display Settings window, 45
 - Settings app, 138
- starting programs from Run window, 49
- Startup and Recovery options (Advanced tab of System applet), 108–110
- Startup section (Apps Settings screen), 141
- Startup tab
 - System Configuration tool, 88, 89, 90
 - Task Manager, 77
- stateless operating systems, 12–13
- states, of windows, 56
- static IP address, 257–258, 511–512, 518–519
- stealth virus, 403
- storage, 18, 301–302
- Storage section (System section), 144
- Storage Spaces, 165
- string, 625–626
- structured data, 667
- Structured Query Language (SQL) injection attack, 424
- styles, Desktop, 38
- subnet mask, 256, 259
- subword tokens, 664
- sudo command, 313, 314, 343
- sunset date, 100
- superuser, 312–313
- supervised learning, 663
- supply chain/pipeline attacks, 425–426
- support, levels of, 753
- surge protectors, 707–708
- surges, 707
- swap files/virtual memory, in order of volatility, 725
- swap partitions, 159
- swipe lock, 521
- synchronizing browser data, 499–500
- synthetic full backups, 779
- sysprep, 21
- System applet, 96, 103–110
- System attribute, 471
- System Components section (System section), 145
- System Configuration tool, 88–90
- system crashes, 649
- System Failure (Advanced tab of System applet), 109
- System File Checker tool, 570–571
- system files and folders, 477–488
 - /System folder, 298–299
- system folders, macOS, 296–299
- system frameworks, 298
- System Information tool, 90–92
- system instability, 553–557
- System library, 297
- system performance, monitoring, 326
- system processes, in order of volatility, 725
- System Protection tab (System applet), 109
- system requirements, 15–20, 176–177
- System Restore, 591–592, 596
- System section, 143–145
- System Settings, 279–288, 648–649

System Startup option (Advanced tab of System applet), 108
 system tray (systray), 46
 system updates/App Store, 302–303
 system uptime, 327–328
 system variables, 107, 626
 SystemRescueCD tool, 164

T

tailgating, 417
 Tamper Protection, 481
 Task Manager, 72–77
 Task Scheduler snap-in, 81
 Taskbar, 45–47, 138
 TCP/IP tab (System Settings screen), 283
 telemetry options, 194
 Temporal Key Integrity Protocol (TKIP), 507
 Terminal, 278
 Terminal Access Controller Access System Plus (TACACS+), 513
 terminology, for operating systems, 8–10
 testing, 407, 414–415, 766
 Text Cursor section (Accessibility Settings screen), 154
 Text Input section
 Display Settings window, 45
 Settings app, 138
 Text Size section (Accessibility Settings screen), 153–154
 Themes section
 Display Settings window, 45
 Settings app, 137
 thermal runaway, 603
 third-party tools, for remote access, 659
 32-bit operating system, 9
 This PC icon, 53
 threat agent, 404
 threats, 418. *See also specific threats*
 ticketing systems, 750–755
 tilde character (!), 329
 time drift, 565–566
 Time & Language screen (Settings app), 128–129
 Time Machine, 288
 time restrictions, setting, 435
 time zone, 98, 196–197, 198
 Time-based One-Time Password (TOTP), 384
 title bar, in windows, 54
 tokenization, 664–665
 tools, 29–31, 277–279, 657–659
 Tools tab (System Configuration tool), 88
 top command, 326–328
 Torvalds, Linus, 11
 traceroute command, 348–349
 tracert command, 224
 traditional software, 25
 transparency, 674
 Trojan horse, 397
 troubleshooting
 about, 542, 614–615
 best practices for malware removal, 589–597
 exam essentials, 615
 exercises, 556–557
 Microsoft Windows OS
 about, 542
 Blue Screen of Death (BSOD), 557–558
 boot problems, 548–552
 degraded performance, 543–548
 frequency shutdowns, 558–559
 low memory warnings, 561–563
 No OS Found, 552–553
 Services Not Starting, 559–561
 slow-loading profiles, 564–565
 steps for, 567–579
 system instability, 553–557
 time drift, 565–566
 USB controller resource warnings, 564
 mobile OS issues, 597–608
 mobile security issues, 608–614
 performance-based question answers, 833
 performance-based questions, 619
 real world scenarios, 561
 review question answers, 815–817
 review questions, 616–618
 security issues
 about, 580–581
 access denied, 584

- browser redirection, 587–588
- browser-related, 585–589
- certificate issues, 586–587
- degraded browser
 - performance, 588–589
- desktop alerts, 582
- disappearing files, 584
- false alerts and hoaxes, 581–582
- network connectivity issues, 581
- OS notifications, 582–583
- OS update failures, 584–585
- permission changes, 584
- pop-ups, 585–586
- renamed system files, 584
- rogue antivirus, 583–584
- Troubleshooting applet, 117–119
- Troubleshooting section (System section), 145
- Trusted Installer, 584
- Trusted Platform Module (TPM), 486
- trusted sources, 524
- two dots (..), 329
- 2.4 GHz, 509
- two-factor (2FA) authentication. *See*
 - multifactor authentication (MFA)

U

- Ultrabooks, 11
- unexpected application behavior, 614
- Unified Extensible Firmware Interface (UEFI),
 - 163, 179, 213, 431, 548–549, 550
- uninterruptible power supply (UPS), 709
- Universal Naming Convention (UNC)
 - path, 250
- Universal Plug and Play (UPnP), 518
- Universal Security Slot (USS), 370–371
- Universal Serial Bus (USB), 115, 214, 372
- Universal Time Code (UTC), 759
- universally unique ID (UUID), 322
- unofficial sources of applications, 609
- unprotected systems, as a vulnerability, 428
- unstructured data, 667
- unsupervised learning, 663
- untrusted sources, 524

- Update Definitions, 482
- updating
 - applications, 529, 603
 - drivers, 305
 - features, 209–212
 - firmware, 305
 - initiating, 646–647
 - operating system (OS), 528–529
 - running, 595–596
 - security, 209–212
- upgrade paths, 34–35
- upgrading
 - firmware, 514
- Windows
 - about, 176–179
 - boot options, 212–216
 - considerations, 216–217
 - editions, 209
 - image deployment, 207
 - process for, 200–205
 - security an feature updates, 209–212
 - side-by-side upgrading Windows
 - 10/11, 206–207
 - zero-touch deployment, 207–209
- U.S. Fire Administration, 700
- USB controller, resource warnings, 564
- USB drives, installation and, 20–21
- USB section (Bluetooth & Devices Settings
 - screen), 148
- USB Selective Suspend, 115
- USB tokens, 384
- use cases, scripting, 640–647
- User Account Control (UAC), 111, 200,
 - 465
- user accounts, Linux, 312–315
- User Accounts applet, 96, 110–111
- user authentication, 243–244, 460–462
- user interaction, 214
- user library, 297–298
- user permissions, restricting, 434–435
- User Profiles settings (Advanced tab of
 - System applet), 107–108
- User State Migration Tool (USMT),
 - 206–207, 208
- user variables, 107, 627

user-initiated installation, 21
 users, 457–460
 /Users folder, 296–297
 Users tab (Task Manager), 77

V

variables, 107, 625–627
 VBScripts, 635–636
 versions, 7–8, 49
 vertical bar (|), 309
 Video Playback section (Apps Settings screen), 141
 Video Random Access Memory (VRAM), requirements for, 18
 video surveillance, 367–368
 video transformers, 666–667
 videoconferencing, 30, 659
 View tab (Folder Options applet), 101, 102
 view type, changing, 59
 virtual memory, 159
 Virtual Network Computing (VNC), 278, 653–654
 Virtual Private Network (VPN), 246–247, 651–653
 Virus & Threat Protection screen (Settings app), 132, 134
 virus transmission in networks, 403–404
 viruses, 399–403
 VirusTotal, 591
 vishing, 415
 Visual Basic scripts, 635–636
 visual editor, 338
 Visual Effects section (Accessibility Settings screen), 154
 voice call, as multifactor authentication (MFA), 383
 Voice Over IP (VoIP), 415
 voice recognition technology, 378
 voltage, current *vs.*, 686
 volume. *See* partition
 volume shadow copy, 773–775
 Volume Snapshot Service (VSS), 773–775
 VPN section (Network & Internet section), 149

W

Waste Electrical and Electronic Equipment (WEEE) Regulations, 717
 web browsers
 redirection of, 587–588
 security
 about, 488
 credentials managers, 493–494
 download and installation, 489–491
 extensions and plug-ins, 491–492
 features, 502–505
 patching, 491
 secure data transfers, 494–495
 settings, 495–502
 websites
 American Customer Satisfaction Index, 790
 computer recycling, 718
 Cybersecurity and Infrastructure Security Agency (CISA), 404
 Google Chrome, 492
 Hiren's BootCD, 442
 iCloud, 526
 Microsoft Download Center, 178
 Microsoft Management Console (MMC), 78
 Nessus vulnerability scanner, 304
 Python, 638
 U.S. Fire Administration, 700
 VirusTotal, 591
 Windows Compatible Products List, 177
 whaling, 417
 while loop, 628
 whoami command, 233–234
 Wi-Fi Protected Access (WPA), 507
 Wi-Fi Protected Access 2 (WPA2), 507–508
 Wi-Fi Protected Access 3 (WPA3), 508
 Wi-Fi Protected Setup (WPS), 512
 wildcards, 59, 221
 windows, 53–56
 Windows
 about, 11, 32
 editions, 33–35
 elements of windows, 53–55
 features, 35–42

- file management, 57–61
- filesystems, 160–161
- interface, 42–53
- rebuilding profiles, 577–579
- resetting, 571
- restoring, 571–574
- states of windows, 56
- Windows 7, 14
- Windows 8, 47
- Windows 10, 15, 16–17, 33–34, 35–42, 176–177
- Windows 11, 14, 16–17, 35–42, 176–177, 213–214
- Windows Backup section (Accounts screen), 152
- Windows batch scripts, 629–631
- Windows Boot Loader, 213
- Windows Boot Manager, 549
- Windows Compatible Products List, 177
- Windows Defender Antivirus, 479–482
- Windows Defender Firewall, 96, 121–122, 252–255, 482–486
- Windows Deployment Services (WDS), 207, 215
- Windows Hello, 42, 463–464
- Windows kernel, 179
- Windows Preinstallation Environment (WinPE), 212
- Windows Recovery Environment (WinRE), 180, 409–410, 553
- Windows Remote Management (WinRM), 657
- Windows Scripting Host (WSH), 636
- Windows Security screen (Settings app), 131–136
- Windows Server Update Services (WSUS), 427, 575
- Windows shell, initializing, 214
- Windows Update (Settings app), 129–131
- Windows Updates, 575–576
- Windows Vista, 14
- Windows XP, 14, 180
- winload.efi, 549
- Winlogon Process, 214
- Winlogon.exxe, 550
- winresume.exe, 549

- WINS tab (System Settings screen), 283
- winver command, 233
- wiping drives, 442
- wire cutters, 699
- wire strippers, 699
- Wired Equivalent Privacy (WEP), 507
- wired networks, 245–246
- wireless access points (WAPs), 604
- wireless MAC filtering, SOHO networks (wireless), 508
- wireless networks, 245–246
- Wireless Wide Area Network (WWAN), 247–248
- word processing tools, 30
- word tokens, 664
- work from home (WFH), 24
- work-arounds, 754
- workgroups, 36, 105, 236–237
- working copies, 301
- working state, 82
- workstation operating systems, 11
- worms, 404
- write-blocker, 725

X

- x64, 9
- x86, 9
- xcopy command, 221–222

Y

- Yellowdog Updater, Modified (YUM) tool, 304, 316
- Your Info section (Accounts screen), 151
- yum command, 318

Z

- zero trust, 379–380
- zero-day attacks, 421
- zero-touch deployment, 207–209

Online Test Bank

To help you study for your CompTIA A+ certification exams, register to gain one year of FREE access after activation to the online interactive test bank—including with your purchase of this book!

To access our learning environment, simply visit www.wiley.com/go/sybextestprep, follow the instructions to register your book, and instantly gain one year of FREE access after activation to:

- All of the chapter questions from this book, so you can practice in a timed and graded setting
- A set of nearly 400 flashcards
- A searchable glossary

Do you need more practice? Check out these resources:

CompTIA A+ Complete Practice Tests: Core 1 Exam 220-1201 and Core 2 Exam 220-1202, Fourth Edition by Audrey O'Shea (ISBN 9781394330331) With hundreds of practice questions for each domain and with additional complete practice exams, the *Practice Tests* series a great way to build your confidence and readiness for exam day.

CompTIA A+ Complete Review Guide: Core 1 Exam 220-1201 and Core 2 Exam 220-1202, Sixth Edition by Troy McMillan (ISBN 9781394330782) The *Review Guide* series is a must-have companion resource to the *CompTIA A+ Complete Study Guide*. It includes practical examples and insights drawn from the real-world experience of PC and hardware technicians, as well as A+ certification exam highlights and concise end-of-chapter reviews.

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.

Your COMPLETE resource for the UPDATED CompTIA A+ Core 1 and Core 2 exams

Building on the popular Sybex self-study approach, this Fifth Edition of bestselling *CompTIA® A+® Complete Study Guide: Core 1 Exam 220-1101 and Core 2 Exam 220-1102*, is revised and updated to cover major advances in technology, including mobile, cloud, networking, and security. The Study Guide includes practical examples and insights drawn from real-world experience, as well as exam highlights and end-of-chapter review. Readers also have access to additional learning resources, including the Sybex interactive online learning environment and test bank with hundreds of review questions, practice exams, electronic flashcards, and a searchable PDF of a Glossary of the most important terms readers will need to understand.

The book provides 100% coverage of the CompTIA A+ Core 1 and Core 2 exam objectives, and verifies the successful candidate has the knowledge and job skills needed for:

- Installing, configuring, and maintaining computer equipment, mobile devices, and software for end users
- Servicing components based on customer requirements
- Understanding networking basics and applying basic cybersecurity methods to mitigate threats
- Properly and safely diagnosing, resolving, and documenting common hardware and software issues
- Applying troubleshooting skills and providing customer support using appropriate communication skills
- Understanding the basics of scripting, cloud technologies, virtualization, and multi-OS deployments in corporate environments

ABOUT CompTIA A+ CERTIFICATION

The CompTIA A+ certification is the industry standard for computer technicians. Exams measure a candidate's competency with various hardware components and software operating systems. A+ certification is the most popular entry-level IT certification, especially for those switching to IT from another career.


ABOUT THE AUTHORS

QUENTIN DOCTER, A+, Network+, IT Fundamentals+, Cloud Essentials+, is an IT consultant who started in the industry in 1994 working in tech support for Packard Bell. Since then, he's worked as a tech and network support specialist, trainer, consultant, and webmaster. During his career, he has achieved certifications from CompTIA, Microsoft, Cisco, Novell, and Sun Microsystems. He has written over a dozen books for Sybex, including the *CompTIA IT Fundamentals+ Study Guide* and the *CompTIA Cloud Essentials+ Study Guide*.

JON BUHAGIAR, A+, Network+, is an information technology professional with two decades of experience in higher education. During the past 22 years he has been responsible for Network Operations at Pittsburgh Technical College and lead several projects, including virtualization (server and desktop), VoIP, Microsoft 365, and many other projects supporting the quality of education at the college. He has achieved several certifications from CompTIA, Cisco, and Microsoft, and taught many of the certification paths. He is the author of several books, including Sybex's *CompTIA Network+ Review Guide*.

www.sybex.com
www.wiley.com/go/sybextestprep

Cover Design: Wiley
Cover Image: © Getty Images Inc./Jeremy Woodhouse




Interactive learning environment

Take your exam prep to the next level with Sybex's superior Interactive online study tools. To access our learning environment, simply visit www.wiley.com/go/sybextestprep, register your book to receive your unique PIN, and instantly gain one year of FREE access after activation to:

- **Interactive test bank with 8 practice exams** to help you identify areas where further review is needed.
- **Over 500 electronic flashcards** to reinforce learning and last-minute prep before the exam
- **Comprehensive glossary** in PDF format gives you instant access to the key terms so you are fully prepared
- **Over 2 hours of recorded audio instruction** from the authors to help you study for the certification exam with confidence

CATEGORY
COMPUTERS/Certification Guides

ISBN 978-1-119-86291-8
56000



9 781119 862918